RIGA: Rotation-Invariant and Globally-Aware Descriptors for Point Cloud Registration

Hao Yu, Ji Hou, Zheng Qin, Mahdi Saleh, Ivan Shugurov, Kai Wang, Benjamin Busam, Slobodan Ilic

Abstract—Successful point cloud registration relies on accurate correspondences established upon powerful descriptors. However, existing neural descriptors either leverage a rotation-variant backbone whose performance declines under large rotations, or encode local geometry that is less distinctive. To address this issue, we introduce RIGA to learn descriptors that are Rotation-Invariant by design and Globally-Aware. From the Point Pair Features (PPFs) of sparse local regions, rotation-invariant local geometry is encoded into geometric descriptors. Global awareness of 3D structures and geometric context is subsequently incorporated, both in a rotation-invariant fashion. More specifically, 3D structures of the whole frame are first represented by our global PPF signatures, from which structural descriptors are learned to help geometric descriptors sense the 3D world beyond local regions. Geometric context from the whole scene is then globally aggregated into descriptors. Finally, the description of sparse regions is interpolated to dense point descriptors, from which correspondences are extracted for registration. To validate our approach, we conduct extensive experiments on both object- and scene-level data. With large rotations, RIGA surpasses the state-of-the-art methods by a margin of $8°$ in terms of the Relative Rotation Error on ModelNet40 and improves the Feature Matching Recall by at least 5 percentage points on 3DLoMatch.

Index Terms—Point Cloud Registration, Rotation-Invariant Descriptors, Globally-Aware Descriptors, Coarse-to-Fine Correspondences

✦

1 INTRODUCTION

OUR entire world is 3D. Modern depth sensors are able to retrieve distance measures of the environment and represent it as point clouds. Naturally, registering point clouds under different sensor poses, a.k.a. point cloud registration, plays a crucial role in a wide range of real applications such as scene reconstruction, autonomous driving, and simultaneous localization and mapping (SLAM). Given a pair of partially-overlapping point clouds, point cloud registration aims to recover the relative transformation between them. As the relative transformation can be solved in closed-form or estimated by a robust estimator [1] based on putative correspondences, establishing reliable correspondences becomes the key to successful registration.

Correspondences are established by matching points according to their associated descriptors. As dense matching is computationally complex, existing works [6], [7], [8], [9], [10], [3], [11], [12], [2], [13] widely adopt a first-sampling-then-matching paradigm to match sparse nodes that are either uniformly-sampled or saliently-detected from dense points. Although the computational complexity is significantly reduced, it introduces a new problem of repeatability, i.e., the corresponding points of some nodes are excluded after sparse sampling s.t. they can never be correctly matched. Due to this design, a considerable part of true correspondences is automatically dropped before matching, which significantly constrains the reliability of putative correspondences. To tackle the problem, we have proposed CoFiNet [14] which extracts hierarchical corre-
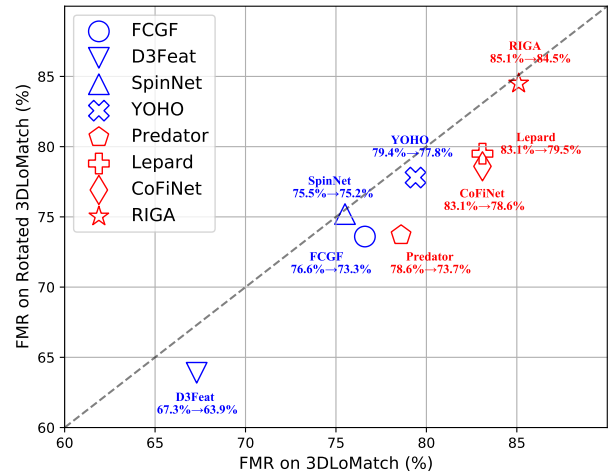

Fig. 1: **Feature Matching Recall (FMR) on 3DLo-Match [2] (x-axis) and Rotated 3DLoMatch (y-axis).** Methods that only encode local geometry are marked as blue, while approaches with global awareness are drawn in red. The performance drop from the original (x-axis) to the rotated (y-axis) benchmark for each method is also demonstrated. Methods that are more robust against rotations are closer to the 45-degree line. Generally, globally-aware methods perform better on standard benchmarks, while rotation-invariant ones (SpinNet [3], YOHO [4] and RIGA) degenerate less under larger rotations. RIGA performs the best in both cases with a drop of only 0.6 percent points.

- *H.Yu, M.Saleh, I.Shugurov, B.Busam and S.Ilic are with Technical University of Munich, Germany. J.Hou is with Meta Reality Labs, USA. Z.Qin is with National University of Defense Technology, China. K.Wang is with Shanghai Jiao Tong University, China. I.Shugurov and S.Ilic are also with Siemens AG, Germany.*
- *E-mail: hao.yu@tum.de*

spondences from coarse to fine. On a coarse scale, it learns to match uniformly-sampled nodes whose vicinities share more overlap. The coarse matching significantly shrinks the space of correspondence search of the consecutive stage,

(a) Point Clouds   **(b) Geometric Descriptors**   **(c) Structural Descriptors**   **(d) Point Descriptors**
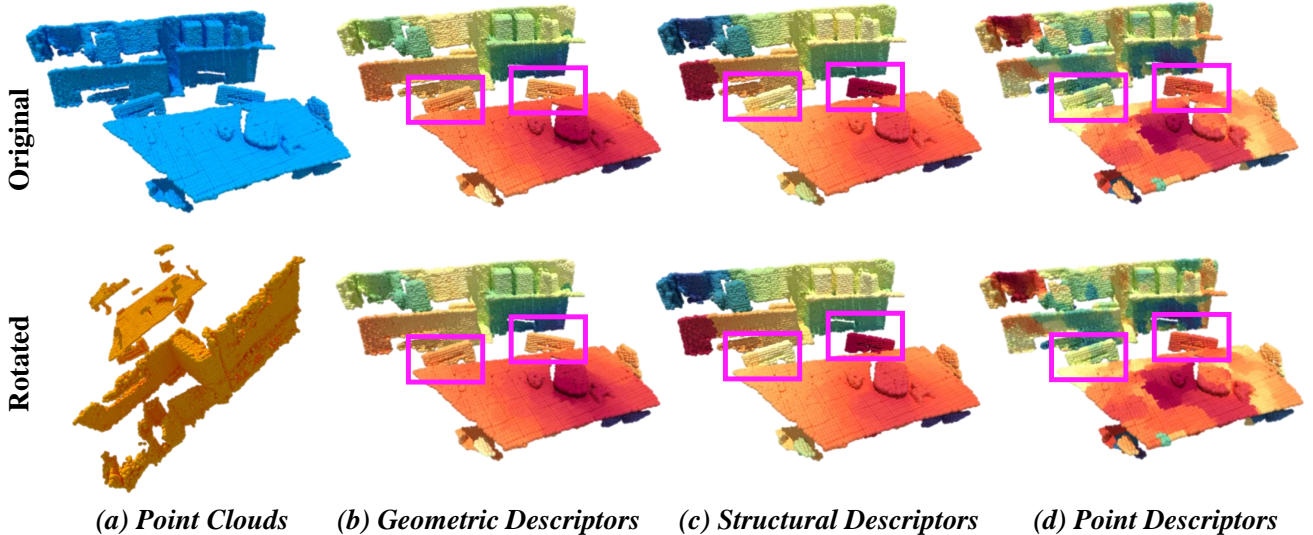
Fig. 2: **Illustration of the Inherent Rotational Invariance and Distinctiveness of RIGA.** In (a), an arbitrary rotation is applied to the input scan. **1) Rotational Invariance**: In (b), (c) and (d), local, global and point descriptors from **untrained** RIGA are visualized by t-SNE [5], respectively. The rotated point cloud is aligned for better visualization. All the descriptors from untrained RIGA remain unchanged after rotation (the second row), which illustrates our inherent rotational invariance guaranteed by design. **2) Distinctiveness**: In (b), although two chairs inside pink rectangles have similar local geometric descriptors, they are distinguishable in (c) where global structures are encoded, and in (d) where global context is incorporated into local descriptors.

where finer correspondences are extracted from the overlapping vicinities. It implicitly considers all the possible correspondences in the matching procedure and therefore eliminates the repeatability issue. However, the descriptors upon which correspondences are extracted by CoFiNet lack robustness against rotations by design. As a consequence, although reliable correspondences are extracted via the proposed coarse-to-fine mechanism, the performance of CoFiNet still significantly declines when rotations are enlarged, as illustrated in Fig. 1.

This phenomenon reminds us of the importance of point descriptors and shifts our attention to introducing more powerful descriptors for better registration performance. Recent trends widely adopt neural backbones [15], [16], [17] to obtain more powerful descriptors [6], [7], [8], [9], [10], [3], [11], [18], [2], [4], [14], [13], [19] from raw points, which gains significant improvement over handcrafted features [20], [21], [22]. The most recent deep learning-based methods [3], [2], [4], [14], [13] can be split into two categories according to the way they enhance descriptors. The first one [3], [4] aims at promising the rotational invariance of descriptors learned from local geometry by design. For a point $\mathbf{x}_i \in \mathbb{R}^3$ from point cloud $\mathcal{X}$, they propose to guarantee that the local descriptor learned from the support area $\Omega_i^{\mathcal{X}}$ around $\mathbf{x}_i$ by a model $\mathcal{G}$ is invariant under arbitrary rotations $\mathbf{R} \in SO(3)$, i.e., $\mathcal{G}(\mathbf{R}(\mathbf{x}_i)|\mathbf{R}(\Omega_i^{\mathcal{X}})) = \mathcal{G}(\mathbf{x}_i|\Omega_i^{\mathcal{X}})$. According to [7], [3], these methods are more robust to larger rotations, which is also demonstrated in Fig. 1 (see SpinNet, YOHO, and RIGA). The second one [2], [14], [13] instead focuses on incorporating global awareness into local descriptors to enhance the distinctiveness. Compared to descriptors that only encode local geometry, i.e., $\mathcal{G}(\mathbf{x}_i|\Omega_i^{\mathcal{X}})$, the globally-aware descriptor $\mathcal{G}(\mathbf{x}_i|\mathcal{X})$ of point $\mathbf{x}_i$ is more distinctive and

much easier to be distinguished from other globally-aware descriptors $\mathcal{G}(\mathbf{x}_j|\mathcal{X})$ of points $\mathbf{x}_j$ with $i \neq j$. As illustrated in Fig. 2, although it is hard to distinguish two chairs according to local geometry (Fig. 2 (b)). However, global awareness helps to separate their description (Fig. 2 (c)). Therefore, globally-aware methods usually perform better on the registration task than approaches that only encode local geometry alone, which is also demonstrated in Fig. 1. However, each category of methods has its specific drawback – rotation-invariant descriptors are usually less distinctive due to the blindness to the global context, while globally-aware methods can produce inconsistent descriptions due to the inherent lack of rotational invariance. The current literature lacks an approach that fulfills both aspects simultaneously, i.e., $\mathcal{G}(\mathbf{R}(\mathbf{x}_i)|\mathbf{R}(\mathcal{X})) = \mathcal{G}(\mathbf{x}_i|\mathcal{X})$.

We propose to bridge the lack of globally-aware descriptors that inherently guarantee rotation invariance for the task of point cloud registration with **RIGA**. Our proposed method simultaneously strengthens the robustness against rotations and distinctiveness of learned descriptors, from which coarse-to-fine keypoint-free correspondences are consecutively extracted. More specifically, we adopt a PointNet [15] architecture, which takes as input the rotation-invariant handcrafted descriptors to encode rotation-invariant local geometry. To provide a node-specific description of the entire scene in a rotation-invariant fashion, we design global PPF signatures that describe each node by considering the spatial relationship of the remaining nodes w.r.t. it. Subsequently, rotation-invariant structural descriptors are learned from global PPF signatures and leveraged to incorporate awareness of global 3D structures into local descriptors. A Transformer [23] architecture is further added, yielding a Vision Transformer (ViT) [24]

architecture to incorporate global awareness of geometric context. Finally, dense point descriptors are obtained by interpolation, and the coarse-to-fine mechanism proposed in CoFiNet [14] is extended to extract reliable correspondences from our rotation-invariant and globally-aware descriptors for point cloud registration.

To the best of our knowledge, RIGA is the first to learn both rotation-invariant and globally-aware descriptors for point cloud registration. Our contributions are summarized as:

- We propose an end-to-end pipeline that guarantees the rotational invariance of globally-aware descriptors by design and extracts coarse-to-fine correspondences for point cloud registration.
- We propose global PPF signatures to provide a node-specific description of the entire scene in a rotation-invariant fashion and further learn global structural descriptors from them to incorporate global structural awareness into local descriptors.
- We empirically show the effectiveness of rotational invariance and global awareness on both object- and scene-level data.

2 RELATED WORK

2.1 Rotation-Invariant Descriptors

2.1.1 Handcrafted Rotation-Invariant Descriptors

Handcrafted rotation-invariant descriptors [20], [21], [22], [25], [26] have been widely explored in 3D by researchers before the popularity of deep neural networks. To guarantee the invariance under rotations, many handcrafted local descriptors [25], [26] rely on an estimated local reference frame (LRF), which is typically based on the covariance analysis of the local surface, to transform local patches to a defined canonical representation. The major drawback of LRF is its non-uniqueness. The constructed rotational invariance is therefore fragile and sensitive to noise. As a result, the attention shifts to those LRF-free approaches [20], [21], [22]. These methods focus on mining the rotation-invariant components of local surfaces and using them to represent the local geometry. Given a point of interest and its adjacent points within the vicinity area, PPF [22] describes each pairwise relationship using Euclidean distances and angles among point vectors and normals. In a similar way, PFH [20] and FPFH [21] encode the geometry of the local surface using the histogram of pairwise geometrical properties. Although these handcrafted descriptors are rotation-invariant by design, all of them are far from satisfactory to be applied in real scenarios with complicated geometry and severe noise.

2.1.2 Learning-based Rotation-Invariant Descriptors

Recently, many deep learning-based methods [7], [8], [3] make the attempt to learn descriptors in a rotation-invariant fashion. As a pioneer, PPF-FoldNet [7] encodes PPF patches into embeddings, from which a FoldingNet [27] decoder reconstructs the input. Correspondences are extracted from the rotation-invariant embeddings for registration. Different from PPF-FoldNet [7] that learns from handcrafted LRF-free descriptors, 3DSN [8] leverages LRF, which transforms local

patches around interest points to defined canonical representations, to enhance the robustness of learned descriptors against rotations. Similarly, SpinNet [3] and Graphite [11], [12] align local patches according to the defined axes before learning descriptors from them. However, all those methods are limited by their locality, i.e., their descriptors are only learned from the local region where their rotational invariance is defined. Those descriptors are blind to the global context and are therefore less distinctive. Without relying on rotation-invariant handcrafted features, YOHO [4] leverages an icosahedral group to learn a group of rotation-equivariant descriptors for each point. Rotating the input point cloud will permute the descriptors within the group, and rotational invariance is achieved by max-pooling over the group. However, its rotational equivariance is fragile in practice, as the finite rotation group cannot span the infinite rotation space. Additionally, expanding a single descriptor to a group damages efficiency. In object-centric registration, recent methods [28], [29], [22] strengthen the rotational invariance in their learned descriptors by concatenating rotation-invariant descriptors, e.g., PPF [22], with their rotation-variant input. However, as shown in Tab. 1, the registration performance of those methods still drops severely when facing large rotations [29].

2.2 Globally-Aware Descriptors

PPF, as an example, has been made semi-global before the widespread of deep neural networks for different tasks [30], [31], [22], [32]. With the widespread of deep neural networks, PPFNet [7] makes the first attempt to incorporate learned global context into their learned descriptors. However, their descriptors are rotation-variant in nature, as the absolute coordinates and PPF features are concatenated as input. Moreover, naively leveraging a max-pooling operator for global awareness largely neglects global information beyond each local patch. Predator [2] leverages attention [23] mechanism in a point cloud registration method to strengthen their descriptors with learned global context. Global information is incorporated from the same and the opposite frame, by interleaving Edge Conv-based [33] self-attention modules and Transformer-based [23] cross-attention modules, respectively. Similarly, Yu et al. [14] interleave Transformer-based [23] self- and cross-attention modules for learning globally-aware descriptors. Such a paradigm is also leveraged in the most recent works [13], [34], [35] for incorporating global awareness into local descriptors. However, these methods ignore the inherent rotational invariance of their learned descriptors. As a result, rotational invariance is learned through data augmentation during training, which is intricate for large rotations and adds significant capacity requirements to the deep model.

3 METHOD

3.1 Problem Statement

We aim at recovering the rigid transformation $\mathbf{T} = \{\mathbf{R} \in \mathcal{SO}(3), \mathbf{t} \in \mathbb{R}^3\}$ that best aligns two partially-overlapping point clouds $\mathcal{X} = \{\mathbf{x}_1, \ldots, \mathbf{x}_N\}$ and $\mathcal{Y} = \{\mathbf{y}_1, \ldots, \mathbf{y}_M\}$. We follow the paradigm of those correspondence-based
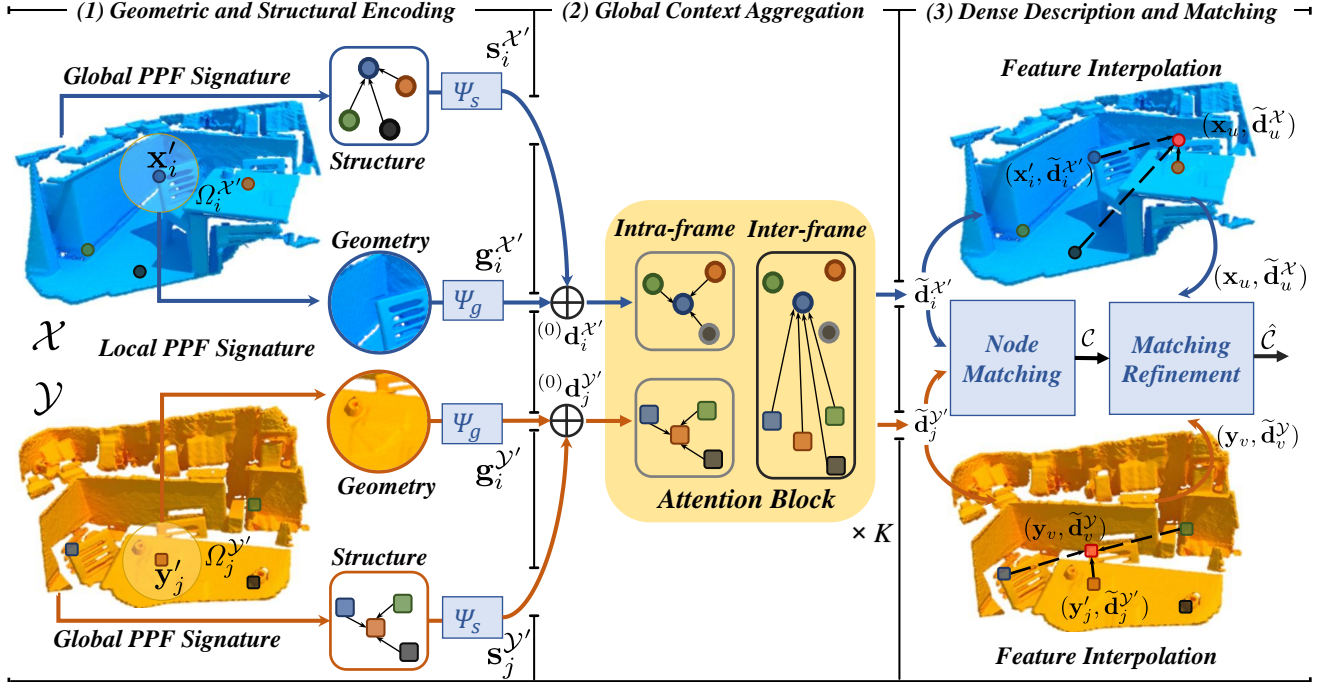
Fig. 3: **Method Overview.** Point cloud $\mathcal{X}$ and $\mathcal{Y}$ are processed in the same way, and we only explain for $\mathcal{X}$ hereafter. **(1)** Local and global PPF signatures are computed for each node $\mathbf{x}_i$, which is sparsely sampled from $\mathcal{X}$. Local geometry and global structures are encoded into descriptors $\mathbf{g}_i^{\mathcal{X}'}$ and $\mathbf{s}_i^{\mathcal{X}'}$ by PointNet [15] Ψ_g and Ψ_s, respectively. **(2)** $\mathbf{s}_i^{\mathcal{X}'}$ joins $\mathbf{g}_i^{\mathcal{X}'}$ with global 3D structures via element-wise addition, yielding a globally-informed descriptor $^{(0)}\mathbf{d}_i^{\mathcal{X}'}$. A stack of K attention blocks is leveraged, where intra- and inter-frame geometric context is globally incorporated, resulting in a globally-aware descriptor $\widetilde{\mathbf{d}}_i^{\mathcal{X}'}$. **(3)** Descriptor $\widetilde{\mathbf{d}}_u^{\mathcal{X}}$ of every point $\mathbf{x}_u \in \mathcal{X}$ is obtained via interpolation. Node correspondence set $\mathcal{C}$ is retrieved in the Node Matching Module (Fig. 5(a)). In the Matching Refinement Module (Fig. 5(b)), point correspondence set $\hat{\mathcal{C}}$ is extracted according to $\mathcal{C}$ and point descriptors. All the descriptors are invariant to rotations by design.

models [6], [9], [10], [3], [2], [14], [35], [34], [13], where transformation is solved based on a putative correspondence set $\hat{\mathcal{C}}$ according to:

$$\arg\min_{\mathbf{R},\mathbf{t}} \sum_{(\mathbf{x}_i,\mathbf{y}_j)\in\hat{\mathcal{C}}} \|\mathbf{R}\cdot\mathbf{x}_i + \mathbf{t} - \mathbf{y}_j\|_2^2, \qquad (1)$$

where $\|\cdot\|_2$ represents the Euclidean norm, and the correspondence set $\hat{\mathcal{C}}$ is established by matching points according to their associated descriptors. In this paper, we focus on learning more powerful descriptors that are inherently rotation-invariant and globally aware. By combining the coarse-to-fine matching mechanism [14], our descriptors lead to more reliable correspondences and thus better registration performance. An overview of the RIGA pipeline can be found in Fig. 3.

3.2 Learning Rotation-Invariant Descriptors from Local Geometry

The first step of our method is the rotation-invariant encoding of geometry within local areas. In the following, we will explain it on the example of $\mathcal{X}$. Encoding is done in exactly the same way for $\mathcal{Y}$. Firstly, N' nodes $\mathcal{X}' = \{\mathbf{x}_1', \mathbf{x}_2', \cdots, \mathbf{x}_{N'}'\}$ are sampled out of N points via Farthest Point Sampling [16]. For each node $\mathbf{x}_i' \in \mathcal{X}'$, its support area $\Omega_i^{\mathcal{X}'}$ can be defined by a radius $r \in \mathbb{R}$, which is demonstrated as:

$$\Omega_i^{\mathcal{X}'} = \{\mathbf{x}_u \in \mathcal{X} | \|\mathbf{x}_i' - \mathbf{x}_u\|_2 < r\}. \qquad (2)$$

Each support area is represented with a set of rotation-invariant PPFs [22]. As shown in Fig. 4(a), for node $\mathbf{x}_i'$, normal $\mathbf{n}_i'$ of $\mathbf{x}_i'$ and $\mathbf{n}_u$ of each point $\mathbf{x}_u \in \Omega_i^{\mathcal{X}'}$ are estimated [36], and the local PPF signature of $\mathbf{x}_i'$ is represented as a set of PPFs:

$$\mathcal{S}_l(\mathbf{x}_i'|\Omega_i^{\mathcal{X}'}) = \{\xi(\mathbf{x}_u, \mathbf{n}_u|\mathbf{x}_i', \mathbf{n}_i')\big|\mathbf{x}_u \in \Omega_i^{\mathcal{X}'}\}, \qquad (3)$$

with each PPF defined as:

$$\xi(\mathbf{x}_u, \mathbf{n}_u|\mathbf{x}_i', \mathbf{n}_i') = (\|\mathbf{d}\|_2, \angle(\mathbf{n}_i', \mathbf{d}), \angle(\mathbf{n}_u, \mathbf{d}), \angle(\mathbf{n}_i', \mathbf{n}_u)), \qquad (4)$$

where $\mathbf{d}$ represents the vector between $\mathbf{x}_i'$ and $\mathbf{x}_u$, and $\angle$ computes the angle between two vectors $\mathbf{v}_1$ and $\mathbf{v}_2$, following the way in [30], [6]:

$$\angle(\mathbf{v}_1, \mathbf{v}_2) = atan2(\|\mathbf{v}_1 \times \mathbf{v}_2\|_2, \mathbf{v}_1 \cdot \mathbf{v}_2). \qquad (5)$$

Then, we leverage PointNet [15] to project each local PPF signature to a c-dimension local geometric descriptor:

$$\mathbf{g}_i^{\mathcal{X}'} = \Psi_g(\mathcal{S}_l(\mathbf{x}_i'|\Omega_i^{\mathcal{X}'})) \in \mathbb{R}^c, \qquad 1 \leq i \leq N', \qquad (6)$$

where Ψ_g stands for a PointNet [15] model shared across all the support areas, and c is the dimension of learned local descriptors. As a result, each support area is described by a rotation-invariant geometric descriptor of length c.

(a) Local PPF Calculation   **(b) Global PPF Calculation**

Fig. 4: **Illustration of PPF Calculation. n** and **n′** denote normals. (a) shows the local PPF of a point $\mathbf{x}_u \in \Omega_i^{\mathcal{X}'}$ with respect to node $\mathbf{x}_i'$. (b) shows the global PPF setup of a node $\mathbf{x}_j'$ sampled from $\mathcal{X}$ with respect to node $\mathbf{x}_i'$.

3.3 Learning Rotation-Invariant Descriptors from Global 3D Structures

The learned geometric descriptor $\mathbf{g}_i^{\mathcal{X}'}$, defined in Eq. 6, is conditioned only on its support area $\Omega_i^{\mathcal{X}'}$. Consequently, it lacks awareness of the global context and is less distinctive for correspondence search. We consider this the main reason why existing rotation-invariant methods [6], [8], [3], [4] fail to compete with rotation-variant but globally-aware approaches [2], [14], [13]. To address this issue, we propose to enrich local descriptors with global structural cues learned from our global PPF signatures that are invariant to rotations by design.

The design of global PPF signatures is inspired by the handcrafted PPF which is widely used for describing local geometry. For each node $\mathbf{x}_i'$ with normal $\mathbf{n}_i'$ ($1 \leq i \leq N'$), we compute the structural relationship of every other node $\mathbf{x}_j' \in \mathcal{X}'$ w.r.t. it (see Fig. 4(b)) by:

$$\mathcal{S}_g(\mathbf{x}_i'|\mathcal{X}') = \{\xi(\mathbf{x}_j', \mathbf{n}_j'|\mathbf{x}_i', \mathbf{n}_i')|\mathbf{x}_j' \in \mathcal{X}', j \neq i\}, \quad (7)$$

which we define as the global PPF signature of node $\mathbf{x}_i'$. Similar to the conventional PPF, the obtained global PPF signatures are rotation-invariant by design. However, the global PPF signatures are unordered as well. Besides, as the global PPF signatures are conditioned on the whole scene represented by sparse nodes, they can be sensitive to partial overlap, i.e., although some nodes can be occluded in $\mathcal{Y}$, they still contribute to the structural awareness of $\mathbf{x}_i'$. Therefore, we further leverage a second PointNet [15] architecture Ψ_s to address both issues simultaneously. The network Ψ_s projects each global PPF signature to a c-dimension structural descriptor. This successfully eliminates the inherent unordered property of the global PPF signatures and provides more robustness against partial overlap in real scenes. We denote the obtained structural descriptors as:

$$\mathbf{s}_i^{\mathcal{X}'} = \Psi_s(\mathcal{S}_g(\mathbf{x}_i'|\mathcal{X}')) \in \mathbb{R}^c, \quad 1 \leq i \leq N'. \quad (8)$$

Each global structural descriptor $\mathbf{s}_i^{\mathcal{X}'}$ will be used to inform its corresponding local geometric descriptor $\mathbf{g}_i^{\mathcal{X}'}$ with global structural information from 3D space.

3.4 Rotation-Invariant Global Awareness

3.4.1 Incorporating Global Information from 3D Structures

Following the examples of [37], [2], [14], we interleave self- and cross-attention for intra- and inter-frame global context,

respectively. However, the standard attention [23] lacks the awareness of global 3D structures, as it is based purely on the similarity of learned geometry. To this end, we inform each learned local geometric descriptor $\mathbf{g}_i^{\mathcal{X}'}$ ($1 \leq i \leq N'$) and $\mathbf{g}_j^{\mathcal{Y}'}$ ($1 \leq j \leq M'$) with global structural cues encoded in corresponding global structural descriptor $\mathbf{s}_i^{\mathcal{X}'}$ and $\mathbf{s}_j^{\mathcal{Y}'}$, respectively. The obtained globally-informed descriptors are calculated as $^{(0)}\mathbf{d}_i^{\mathcal{X}'} = \mathbf{g}_i^{\mathcal{X}'} \oplus \mathbf{s}_i^{\mathcal{X}'}$ and $^{(0)}\mathbf{d}_j^{\mathcal{Y}'} = \mathbf{g}_j^{\mathcal{Y}'} \oplus \mathbf{s}_j^{\mathcal{Y}'}$, where $\oplus$ is the element-wise addition.

3.4.2 Global Intra-Frame Aggregation of Geometric Context

A stack of K attention blocks operates on globally-informed descriptors to exchange learned geometric information among nodes. Each attention block has an intra-frame module followed by an inter-frame module.

Taking node $\mathbf{x}_i' \in \mathcal{X}'$ as an example, we detail the computation of the intra-frame module inside the l^{th} ($1 \leq l \leq K$) attention block hereafter. Learnable matrices $^{(l)}\mathbf{W}_q$, $^{(l)}\mathbf{W}_k$, and $^{(l)}\mathbf{W}_v \in \mathbb{R}^{c \times c}$ are introduced to linearly project $^{(l-1)}\mathbf{d}_i^{\mathcal{X}'}$ to *query*, *key*, and *value* with:

$$
\begin{aligned}
^{(l)}\mathbf{q}_i^{\mathcal{X}'} &= {}^{(l)}\mathbf{W}_q \cdot {}^{(l-1)}\mathbf{d}_i^{\mathcal{X}'}, \\
^{(l)}\mathbf{k}_i^{\mathcal{X}'} &= {}^{(l)}\mathbf{W}_k \cdot {}^{(l-1)}\mathbf{d}_i^{\mathcal{X}'}, \\
^{(l)}\mathbf{v}_i^{\mathcal{X}'} &= {}^{(l)}\mathbf{W}_v \cdot {}^{(l-1)}\mathbf{d}_i^{\mathcal{X}'},
\end{aligned}
\quad (9)
$$

respectively, where $^{(l)}\mathbf{q}_i^{\mathcal{X}'}$ and $^{(l)}\mathbf{k}_i^{\mathcal{X}'}$ are used for retrieving similar nodes, and $^{(l)}\mathbf{v}_i^{\mathcal{X}'}$ encodes the context for aggregation.

The attention [23] is defined on a node set $\mathcal{S} \in \{\mathcal{X}', \mathcal{Y}'\}$:

$$^{(l)}\mathbf{a}_i^{\mathcal{X}' \leftarrow \mathcal{S}} = \text{softmax}([^{(l)}a_i^1, {}^{(l)}a_i^2, \cdots, {}^{(l)}a_i^{|\mathcal{S}|}])^T / \sqrt{c} \in \mathbb{R}^{|\mathcal{S}|}, \quad (10)$$

where $^{(l)}a_i^j$ is calculated as $^{(l)}a_i^j = (^{(l)}\mathbf{q}_i^{\mathcal{X}'})^T \cdot {}^{(l)}\mathbf{k}_j^{\mathcal{S}}$ ($1 \leq j \leq |\mathcal{S}|$), and $|\cdot|$ denotes the set cardinality. The message $^{(l)}\mathbf{m}_i^{\mathcal{X}' \leftarrow \mathcal{S}} \in \mathbb{R}^c$, which flows from set $\mathcal{S}$ to node $\mathbf{x}_i' \in \mathcal{X}'$, is calculated as:

$$^{(l)}\mathbf{m}_i^{\mathcal{X}' \leftarrow \mathcal{S}} = [^{(l)}\mathbf{v}_1^{\mathcal{S}}, {}^{(l)}\mathbf{v}_2^{\mathcal{S}}, \cdots, {}^{(l)}\mathbf{v}_{|\mathcal{S}|}^{\mathcal{S}}] \cdot {}^{(l)}\mathbf{a}_i^{\mathcal{X}' \leftarrow \mathcal{S}} \in \mathbb{R}^c. \quad (11)$$

We globally aggregate the intra-frame learned geometry with:

$$^{(l)}\overline{\mathbf{d}}_i^{\mathcal{X}'} = {}^{(l-1)}\mathbf{d}_i^{\mathcal{X}'} + \text{MLP}([^{(l-1)}\mathbf{d}_i^{\mathcal{X}'}, \mathbf{m}_i^{\mathcal{X}' \leftarrow \mathcal{S}}]), \quad (12)$$

where MLP is a multilayer perceptron with $\mathcal{S} = \mathcal{X}'$. For node $\mathbf{y}_j' \in \mathcal{Y}'$, $^{(l)}\overline{\mathbf{d}}_j^{\mathcal{Y}'}$ is calculated in the same way according to Eq. 12, but with $\mathcal{S} = \mathcal{Y}'$.

3.4.3 Global Inter-Frame Fusion of Geometric Context

For the l^{th} ($1 \leq l \leq K$) attention block, the inter-frame module takes as input the output of the intra-frame module, i.e., $^{(l)}\overline{\mathbf{d}}_i^{\mathcal{X}'}$ and $^{(l)}\overline{\mathbf{d}}_j^{\mathcal{Y}'}$. Taking node $\mathbf{x}_i' \in \mathcal{X}'$ as an example, similar to Eq. 9, $^{(l)}\overline{\mathbf{d}}_i^{\mathcal{X}'}$ is linearly projected by learnable matrices $^{(l)}\overline{\mathbf{W}}_q$, $^{(l)}\overline{\mathbf{W}}_k$, and $^{(l)}\overline{\mathbf{W}}_v \in \mathbb{R}^{c \times c}$:

$$
\begin{aligned}
^{(l)}\overline{\mathbf{q}}_i^{\mathcal{X}'} &= {}^{(l)}\overline{\mathbf{W}}_q \cdot {}^{(l)}\overline{\mathbf{d}}_i^{\mathcal{X}'}, \\
^{(l)}\overline{\mathbf{k}}_i^{\mathcal{X}'} &= {}^{(l)}\overline{\mathbf{W}}_k \cdot {}^{(l)}\overline{\mathbf{d}}_i^{\mathcal{X}'}, \\
^{(l)}\overline{\mathbf{v}}_i^{\mathcal{X}'} &= {}^{(l)}\overline{\mathbf{W}}_v \cdot {}^{(l)}\overline{\mathbf{d}}_i^{\mathcal{X}'},
\end{aligned}
\quad (13)
$$

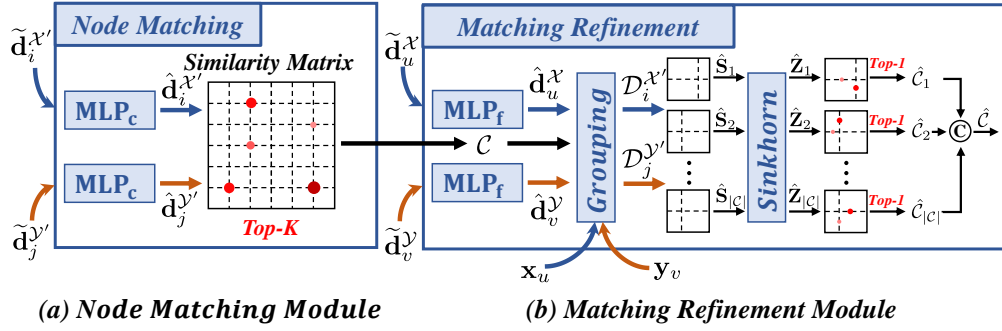*(a) Node Matching Module*   *(b) Matching Refinement Module*

Fig. 5: **Illustration of Coarse-to-Fine Correspondence Extraction.** In (a), nodes from two frames are matched according to the similarity of the MLP-projected descriptors, and node correspondences with Top-*K* highest scores are selected. In (b), according to Eq. 16, each node is assigned with a group of neighbor points, together with their associated MLP-projected descriptors. For each node correspondence, the similarity between their neighbor points is computed. The resulting similarity matrix is normalized by Sinkhorn [38] algorithm. A point correspondence set is extracted from each normalized matrix, and the final point correspondence set is constructed as the union of all the individual ones.

upon which $^{(l)}\overline{\mathbf{a}}_i^{\mathcal{X}'\leftarrow\mathcal{S}}$ and $^{(l)}\overline{\mathbf{m}}_i^{\mathcal{X}'\leftarrow\mathcal{S}}$ are computed following Eq. 10 and Eq. 11, respectively, with $\mathcal{S} = \mathcal{Y}'$. Finally, the geometric context from the opposite frame, i.e., the node set $\mathcal{Y}'$, is fused to node $\mathbf{x}_i'$:

$$^{(l)}\mathbf{d}_i^{\mathcal{X}'} = {}^{(l)}\overline{\mathbf{d}}_i^{\mathcal{X}'} + \text{MLP}([{}^{(l)}\overline{\mathbf{d}}_i^{\mathcal{X}'}, \overline{\mathbf{m}}_i^{\mathcal{X}'\leftarrow\mathcal{S}}]), \qquad (14)$$

with $\mathcal{S} = \mathcal{Y}'$. For node $\mathbf{y}_j' \in \mathcal{Y}'$, $^{(l)}\mathbf{d}_j^{\mathcal{Y}'}$ is calculated in the same way according to Eq. 14, but with $\mathcal{S} = \mathcal{X}'$.

Since all the operation is performed in feature space, the rotation-invariance of $^{(0)}\mathbf{d}_i^{\mathcal{X}'}$ remains in all $^{(l)}\overline{\mathbf{d}}_i^{\mathcal{X}'}$ and $^{(l)}\mathbf{d}_i^{\mathcal{X}'}$ with $1 \leq l \leq K$. As a result, the obtained globally-aware descriptor $\widetilde{\mathbf{d}}_i^{\mathcal{X}'} := {}^{(K)}\mathbf{d}_i^{\mathcal{X}'}$ is rotation-invariant by design. Similarly, globally-aware descriptor $\widetilde{\mathbf{d}}_j^{\mathcal{Y}'} := {}^{(K)}\mathbf{d}_j^{\mathcal{Y}'}$ is also rotation-invariant for each $\mathbf{y}_j' \in \mathcal{Y}$.

3.5 Rotation-Invariant Dense Description

Until here, we have successfully incorporated global awareness into learned local descriptors of nodes without sacrificing the inherent rotational invariance. The aforementioned repeatability issue of sparsely sampled nodes, however, still remains. To address this issue, we leverage the coarse-to-fine strategy proposed in [14], where nodes are first matched according to the overlap ratios of their vicinities, and point correspondences are then extracted from the vicinities of matched nodes. As the first step, dense point descriptors are generated via interpolation. For each point $\mathbf{x}_u \in \mathcal{X}$, we find its *k*-nearest neighbor nodes in $\mathcal{X}'$ according to their Euclidean distance. The descriptor $\widetilde{\mathbf{d}}_u^{\mathcal{X}}$ of point $\mathbf{x}_u$ can be interpolated as:

$$\widetilde{\mathbf{d}}_u^{\mathcal{X}} = \sum_{i=1}^{k} w_i^u \cdot \widetilde{\mathbf{d}}_i^{\mathcal{X}'}, \qquad \text{with} \quad w_i^u = \frac{1/d_i^u}{\sum_{l=1}^{k} 1/d_l^u}, \qquad (15)$$

where d_l^u depicts the Euclidean distance of point $\mathbf{x}_u$ to its l^{th} nearest node in geometry space. Point descriptor $\widetilde{\mathbf{d}}_v^{\mathcal{Y}}$ of $\mathbf{y}_v \in \mathcal{Y}$ is calculated in the same way. As the interpolation coefficients are only related to Euclidean distance, the obtained point descriptors remain invariant to rotations.

3.6 Coarse-to-Fine Correspondence Extraction

The coarse-to-fine mechanism [14] is leveraged to extract correspondences from our obtained node and point descriptors. We first project $\widetilde{\mathbf{d}}_i^{\mathcal{X}'}$ and $\widetilde{\mathbf{d}}_u^{\mathcal{X}}$ by using two individual multilayer perceptrons (MLP), which provides $\hat{\mathbf{d}}_i^{\mathcal{X}'}$ and $\hat{\mathbf{d}}_u^{\mathcal{X}}$ in Fig. 5(a) and (b), respectively. We also project descriptors from point cloud $\mathcal{Y}$ to $\hat{\mathbf{d}}_j^{\mathcal{Y}'}$ and $\hat{\mathbf{d}}_v^{\mathcal{Y}}$. On the coarse level, as shown in Fig. 5(a), the similarity between node $\mathbf{x}_i' \in \mathcal{X}'$ and $\mathbf{y}_j' \in \mathcal{Y}'$ is calculated as $1/\|\hat{\mathbf{d}}_i^{\mathcal{X}'} - \hat{\mathbf{d}}_j^{\mathcal{Y}'}\|_2$. As the following step, Top-*K* node correspondences with the highest similarity values are sampled, resulting in the node correspondence set $\mathcal{C}$ with $|\mathcal{C}|$ correspondences. In "Grouping" of Fig. 5(b), vicinities $(\mathcal{V}_i^{\mathcal{X}'}, \mathcal{V}_j^{\mathcal{Y}'})$ of coarse correspondence $C_l := (\mathbf{x}_i', \mathbf{y}_j') \in \mathcal{C}$ are collected by the point-to-node assignment [39], [14], i.e., assigning points to their nearest nodes in geometry space. For node $\mathbf{x}_i'$, its vicinity $\mathcal{V}_i^{\mathcal{X}'}$ and the associated descriptor group $\mathcal{D}_i^{\mathcal{X}'}$ can be defined as:

$$\begin{cases} \mathcal{V}_i^{\mathcal{X}'} = \{\mathbf{x}_u \in \mathcal{X} | \|\mathbf{x}_u - \mathbf{x}_i'\|_2 < \|\mathbf{x}_u - \mathbf{x}_j'\|_2, \forall j \neq i\}, \\ \mathcal{D}_i^{\mathcal{X}'} = \{\hat{\mathbf{d}}_u^{\mathcal{X}} | \hat{\mathbf{d}}_u^{\mathcal{X}} \leftrightarrow \mathbf{x}_u \ with \ \mathbf{x}_u \in \mathcal{V}_i^{\mathcal{X}'}\}, \end{cases}$$

$$(16)$$

where $\hat{\mathbf{d}}_u^{\mathcal{X}} \leftrightarrow \mathbf{x}_u$ denotes that $\hat{\mathbf{d}}_u^{\mathcal{X}}$ is the descriptor associated to point $\mathbf{x}_u$. $\mathcal{V}_j^{\mathcal{Y}'}$ and $\mathcal{D}_j^{\mathcal{Y}'}$ are defined in the same way for nodes $\mathbf{y}_j' \in \mathcal{Y}'$. Finally, we present the similarity of $(\mathcal{D}_i^{\mathcal{X}'}, \mathcal{D}_j^{\mathcal{Y}'})$ as a matrix $\hat{\mathbf{S}}_l \in \mathbb{R}^{|\mathcal{D}_i^{\mathcal{X}'}| \times |\mathcal{D}_j^{\mathcal{Y}'}|}$, where each entry is calculated as $\hat{\mathbf{S}}_l^{u,v} = (\hat{\mathbf{d}}_u^{\mathcal{X}})^T \cdot \hat{\mathbf{d}}_v^{\mathcal{Y}}$, with $\hat{\mathbf{d}}_u^{\mathcal{X}} \in \mathcal{D}_i^{\mathcal{X}'}$ and $\hat{\mathbf{d}}_v^{\mathcal{Y}} \in \mathcal{D}_j^{\mathcal{Y}'}$. To deal with partial overlap, we follow the slack idea [37] and augment $\hat{\mathbf{S}}_l$ with an additional row and an additional column filled with the same learnable parameter α. In "Sinkhorn" of Fig. 5(b), each augmented similarity matrix is normalized to a confidence matrix $\hat{\mathbf{Z}}_l \in \mathbb{R}^{|\mathcal{D}_i^{\mathcal{X}'}|+1| \times |\mathcal{D}_j^{\mathcal{Y}'}|+1|}$, which is a non-negative matrix with every row and every column summing to 1, with the Sinkhorn [38] algorithm. From $\hat{\mathbf{Z}}_l$ we extract the point correspondence set $\hat{\mathcal{C}}_l$ as the maximum confidence individually for each row and column. The union of all $\hat{\mathcal{C}}_l$ ($1 \leq l \leq |\mathcal{C}|$) constructs the final point correspondence set $\hat{\mathcal{C}}$, which we use for registration.
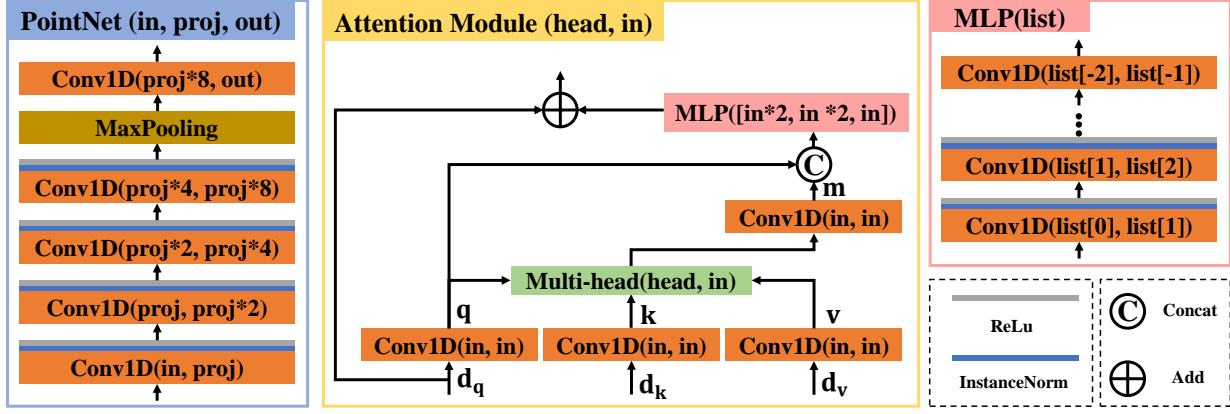
Fig. 6: **Detailed Architecture of Components**. In attention modules, "Multi-head" stands for the multi-head mechanism [23], where $\mathbf{q}$, $\mathbf{k}$ and $\mathbf{v} \in \mathbb{R}^{\mathbf{in}}$ are first reshaped to $(\mathbf{head}, \mathbf{in}/\mathbf{head})$, and attention is then computed separately for each head channel from corresponding $\mathbf{q}$ and $\mathbf{k}$. *Value* $\mathbf{v}$ in each head channel is fused independently according to the attention computed for the same head. The fused *values* with shape $(\mathbf{head}, \mathbf{in}/\mathbf{head})$ are reshaped back to $(\mathbf{in}, \mathbf{1})$, which is finally projected to message $\mathbf{m} \in \mathbb{R}^{\mathbf{in}}$.

3.7 Loss Functions

The total loss function $\mathcal{L} = \mathcal{L}_c + \lambda \mathcal{L}_f$ consists of a coarse-level matching loss $\mathcal{L}_c$ and a fine-scale correspondence refinement loss $\mathcal{L}_f$. $\lambda \in \mathbb{R}$ is the hyper-parameter used to balance the two terms.

3.7.1 Coarse-level Loss for Node Matching

Following [14], our coarse-level loss is defined according to the overlap ratios of the vicinities $(\mathcal{V}_i^{\mathcal{X}'}, \mathcal{V}_j^{\mathcal{Y}'})$ of each node correspondence $(\mathbf{x}_i', \mathbf{y}_j')$. Given vicinities $(\mathcal{V}_i^{\mathcal{X}'}, \mathcal{V}_j^{\mathcal{Y}'})$ of node correspondence $(\mathbf{x}_i', \mathbf{y}_j')$, the number of visible points in one vicinity w.r.t. the other vicinity is defined as:

$$n_i^j = \sum_{\mathbf{x}_u \in \mathcal{V}_i^{\mathcal{X}'}} \mathbb{1}(\exists \mathbf{y}_v \in \mathcal{V}_j^{\mathcal{Y}'} s.t. \|\mathbf{T}(\mathbf{x}_u) - \mathbf{y}_v\|_2 < \tau_p), \quad (17)$$

and

$$n_j^i = \sum_{\mathbf{y}_v \in \mathcal{V}_j^{\mathcal{Y}'}} \mathbb{1}(\exists \mathbf{x}_u \in \mathcal{V}_i^{\mathcal{X}'} s.t. \|\mathbf{T}(\mathbf{x}_u) - \mathbf{y}_v\|_2 < \tau_p), \quad (18)$$

for vicinities $\mathcal{V}_i^{\mathcal{X}'}$ and $\mathcal{V}_j^{\mathcal{Y}'}$, respectively, where $\tau_p \in \mathbb{R}$ is the distance threshold for correspondence decision. The overlap ratio between vicinities $(\mathcal{V}_i^{\mathcal{X}'}, \mathcal{V}_j^{\mathcal{Y}'})$ is further defined as $r_i^j = \frac{1}{2}(\frac{n_i^j}{|\mathcal{V}_i^{\mathcal{X}'}|} + \frac{n_j^i}{|\mathcal{V}_j^{\mathcal{Y}'}|})$.

Similar to [10], [2], [35], we use Circle Loss [40], a variant of Triplet Loss [41], to guide the learning of node descriptors. For a node $\mathbf{x}_i'$ from $\mathcal{X}'$, we sample a positive set $\mathcal{E}_p^i$ composed of nodes $\mathbf{y}_j'$ from $\mathcal{Y}'$ s.t. $\mathbf{T}(\mathcal{V}_i^{\mathcal{X}'})$ overlaps with $\mathcal{V}_j^{\mathcal{Y}'}$, and a negative set $\mathcal{E}_n^i$ consisting of nodes $\mathbf{y}_l'$ from $\mathcal{Y}'$ s.t. $\mathbf{T}(\mathcal{V}_i^{\mathcal{X}'})$ and $\mathcal{V}_l^{\mathcal{Y}'}$ share no overlap, where $\mathbf{T}(\mathcal{V}_i^{\mathcal{X}'})$ denote $\mathcal{V}_i^{\mathcal{X}'}$ transformed by the ground truth transformation $\mathbf{T}$. The loss function on $\mathcal{X}'$ can be defined upon n nodes $\mathbf{x}_i'$ sampled from $\mathcal{X}'$ as:

$$\mathcal{L}_c^{\mathcal{X}'} = \frac{1}{n} \sum_{i=1}^{n} \log \Big[1 + \sum_{\mathbf{y}_j' \in \mathcal{E}_p^i} e^{r_i^j \beta_p^j (d_i^j - \Delta_p)} \cdot \sum_{\mathbf{y}_l' \in \mathcal{E}_n^i} e^{\beta_n^l (\Delta_n - d_i^l)} \Big], \quad (19)$$

where r_i^j is the overlap ratio between $\mathcal{V}_i^{\mathcal{X}'}$ and $\mathcal{V}_j^{\mathcal{Y}'}$, and $d_i^j = \|\hat{\mathbf{d}}_i^{\mathcal{X}'} - \hat{\mathbf{d}}_j^{\mathcal{Y}'}\|_2$ denotes the Euclidean distance of nodes $\mathbf{x}_i'$ and $\mathbf{y}_j'$ in learned feature space. Δ_p and Δ_n are the positive and negative margins, which are set to 0.1 and 1.4 in practice, respectively. Furthermore, $\beta_p^j = \gamma(d_i^j - \Delta_p)$ and $\beta_n^l = \gamma(\Delta_n - d_i^l)$ are the weights determined for each sample individually, with the same hyper-parameter $\gamma \in \mathbb{R}$. We can similarly define the loss $\mathcal{L}_c^{\mathcal{Y}'}$ and write the total coarse-level loss as $\mathcal{L}_c = \frac{1}{2}(\mathcal{L}_c^{\mathcal{X}'} + \mathcal{L}_c^{\mathcal{Y}'})$.

3.7.2 Fine-level Loss for Correspondence Refinement

After getting the coarse correspondence set $\mathcal{C}$, we adopt a negative log-likelihood loss [37] to guide the correspondence refinement procedure. For node correspondence $C_l := (\mathbf{x}_i', \mathbf{y}_j') \in \mathcal{C}$, as mentioned before, we compute its confidence matrix $\hat{\mathbf{Z}}_l \in \mathbb{R}^{|\mathcal{D}_i^{\mathcal{X}'}+1| \times |\mathcal{D}_j^{\mathcal{Y}'}+1|}$ augmented with a slack row and slack column for no correspondence. The ground truth point correspondence set between vicinities $\mathcal{V}_i^{\mathcal{X}'}$ and $\mathcal{V}_j^{\mathcal{Y}'}$ is denoted as $\mathcal{M}_l$, while the sets of unmatched points in vicinity $\mathcal{V}_i^{\mathcal{X}'}$ and $\mathcal{V}_j^{\mathcal{Y}'}$ are represented as $\mathcal{I}_l$ and $\mathcal{J}_l$, respectively. The ground truth point correspondence set between vicinities $\mathcal{V}_i^{\mathcal{X}'}$ and $\mathcal{V}_j^{\mathcal{Y}'}$ is defined as:

$$\mathcal{M}_l = \{(\mathbf{x}_u \in \mathcal{V}_i^{\mathcal{X}'}, \mathbf{y}_v \in \mathcal{V}_j^{\mathcal{Y}'}) | \|\mathbf{T}(\mathbf{x}_u) - \mathbf{y}_v\|_2 < \tau_p\}. \quad (20)$$

The set of occluded points in one vicinity w.r.t. the other one is defined as:

$$\mathcal{I}_l = \{\mathbf{x}_u \in \mathcal{V}_i^{\mathcal{X}'} | \nexists \mathbf{y}_v \in \mathcal{V}_j^{\mathcal{Y}'} s.t. \|\mathbf{T}(\mathbf{x}_u) - \mathbf{y}_v\|_2 < \tau_p\}, \quad (21)$$

and

$$\mathcal{J}_l = \{\mathbf{y}_v \in \mathcal{V}_j^{\mathcal{Y}'} | \nexists \mathbf{x}_u \in \mathcal{V}_i^{\mathcal{X}'} s.t. \|\mathbf{T}(\mathbf{x}_u) - \mathbf{y}_v\|_2 < \tau_p\}, \quad (22)$$

for vicinities $\mathcal{V}_i^{\mathcal{X}'}$ and $\mathcal{V}_j^{\mathcal{Y}'}$, respectively.

Finally, the correspondence refinement loss of C_l reads as:

$$\mathcal{L}_f^l = - \sum_{(\mathbf{x}_u,\mathbf{y}_v) \in \mathcal{M}_l} \log \hat{\mathbf{Z}}_l^{u,v} - \sum_{\mathbf{x}_u \in \mathcal{I}_l} \log \hat{\mathbf{Z}}_l^{u,|\mathcal{D}_j^{\mathcal{Y}'}|+1}$$
$$- \sum_{\mathbf{y}_v \in \mathcal{J}_l} \log \hat{\mathbf{Z}}_l^{|\mathcal{D}_i^{\mathcal{X}'}|+1,v} \tag{23}$$

where $\hat{\mathbf{Z}}_l^{u,v}$ denotes the entry of $\hat{\mathbf{Z}}_l$ on the u^{th} row and v^{th} column. The total loss is averaged across the whole node correspondence set $\mathcal{C}$ as $\mathcal{L}_f = \frac{1}{|\mathcal{C}|} \sum_{l=0}^{|\mathcal{C}|} \mathcal{L}_f^l$.

4 RESULTS

We evaluate RIGA on both synthetic object dataset ModelNet40 [42] and real scene benchmarks, including 3DMatch [43] and 3DLoMatch [2]. RANSAC [1] is leveraged to estimate transformation based on putative correspondences. We further demonstrate our robustness against poor normal estimation in the Appendix by using KITTI [44]. We also compare RIGA to the state-of-the-art methods in terms of inference speed in the Appendix. Qualitative results can be found in Fig. 7. We also illustrate failed cases from 3DLoMatch in Fig. 8. More qualitative results on ModelNet40, 3DMatch, and 3DLoMatch are provided in the Appendix.

4.1 Implementation Details

4.1.1 Detailed Architecture

The detailed architecture of each component leveraged in RIGA can be found in Fig. 6. PointNets [15] Ψ_g and Ψ_s are two individual models with the same architecture (input dimension $\mathbf{in} = 4$, project dimension $\mathbf{proj} = 64$ and output dimension $\mathbf{out} = 256$), as shown in the leftmost column in Fig. 6. Each attention block has an intra-frame module and an inter-frame module, both with the architecture of the "Attention Module" shown in Fig. 6. Differently, for intra-frame modules, $\mathbf{d}_q$, $\mathbf{d}_k$ and $\mathbf{d}_v$ are all from the same frame, while in inter-frame modules, $\mathbf{d}_k$ and $\mathbf{d}_v$ are from the opposite frame. $\mathbf{MLP_c}$ and $\mathbf{MLP_f}$ in Fig. 5 have the same MLP architecture shown in the rightmost column of Fig. 6, with a input dimension list of [256, 128, 64, 32].

4.1.2 Training and Testing

RIGA is implemented with PyTorch [45] and trained end-to-end on a single NVIDIA RTX 3090 with 24G memory, where the batch size is set to 2 for 3DMatch/3DLoMatch [43], [2] and 16 for ModelNet40 [42]. Notably, it could also be trained on a GPU with 11G memory, e.g., NVIDIA GTX 1080Ti. We train for 150 epochs on ModelNet40 and for 20 epochs on 3DMatch/3DLoMatch, both with $\lambda = 1$ to balance different loss functions. We leverage an Adam optimizer [46] with an initial learning rate of 1e-4, which is exponentially decayed by 0.05 after each epoch. On ModelNet40, we sparsely sample $N' = M' = 256$ nodes from each point cloud pair, with a radius $r = 0.2m$ to construct support areas, within which the number of points is truncated to 64. On 3DMatch/3DLoMatch, N' and M' are both set to 512, with $r = 0.3m$ and 512 points within each support area. Besides, the number of points in vicinity $\mathcal{V}$ is truncated to 32 and 128 on ModelNet40 and 3DMatch/3DLoMatch respectively. On both datasets, the dimension of intermediate descriptors

$\mathbf{g}$, $\mathbf{s}$ and $\widetilde{\mathbf{d}}$ is set to 256, while that of descriptors $\hat{\mathbf{d}}$, from which correspondences are hierarchically extracted, is set to 32. The number of neighbor points used for feature interpolation is set to $k = 3$. We use 100 iterations for Sinkhorn [38] algorithm. The number of attention blocks is set to K=6, and the attention mechanism is implemented with 4 heads. During training, 256 node pairs that overlap under ground truth transformation are sampled as the node correspondence set $\mathcal{C}$. During testing, 256 node correspondences with the highest similarity scores are selected for the consecutive refinement.

4.2 Synthetic Object Dataset: ModelNet40

4.2.1 Dataset

ModelNet40 [42] consists of 12,311 CAD models of objects from 40 different categories. We follow the setting of [47], where 9,833 shapes are used for training, and the rest 2,468 for testing. For each model, 1,024 points are randomly sampled from its surface. For simulating the partial overlap from scanning, 768 points nearest to a randomly selected viewpoint in the space are resampled from the 1,024 points, which serves as the input point cloud. Following [29], instead of using the ground truth normals, we estimate them using Open3D [48].

4.2.2 Metrics

We use 3 widely-adopted metrics [29]: (1) *Relative Rotation Error* (**RRE**) that evaluates the error between estimated and ground truth rotation matrices; (2) *Relative Translation Error* (**RTE**) that measures the error between estimated and ground truth translation vectors; (3) *Root-Mean-Square Error* (**RMSE**) which calculates the residual error between correspondences from the same point cloud, separately transformed by the estimated and ground truth transformation. Please refer to the Appendix for the detailed definition.

4.2.3 Comparisons to the State-of-the-Art

We compare RIGA with 9 state-of-the-art baselines, including 7 direct registration methods and 2 correspondence-based approaches (Predator [2] and CoFiNet [14]). The detailed results are shown in Tab. 1. From the second column that lists the dimension of descriptors used for correspondence search, it can be noticed that RIGA uses the most compact descriptors among all the methods. On the "Unseen" setting, RIGA surpasses all the other methods with rotations in the range of [0, 45°]. With a maximum rotation of 180°, it achieves on-par performance with GMCNet [29] and outperforms others. When Gaussian noise is added, although RIGA stays comparable with GMCNet [29] with rotations in [0, 45°], it outperforms all the baselines on all the metrics by a large margin with rotations enlarged to 180°. Notably, all the methods except for RIGA degenerate significantly, which shows the superiority of the inherent rotational invariance of RIGA. Although direct registration methods are specifically tuned with good performance on object-level data as pointed out in [2], RIGA could compete with them and even performs significantly better than them on data with Gaussian noise and large rotations. Moreover, RIGA also achieves the state-of-the-art performance on scene-level benchmarks [43], [2], while most direct registration methods fail to work there according to [2].
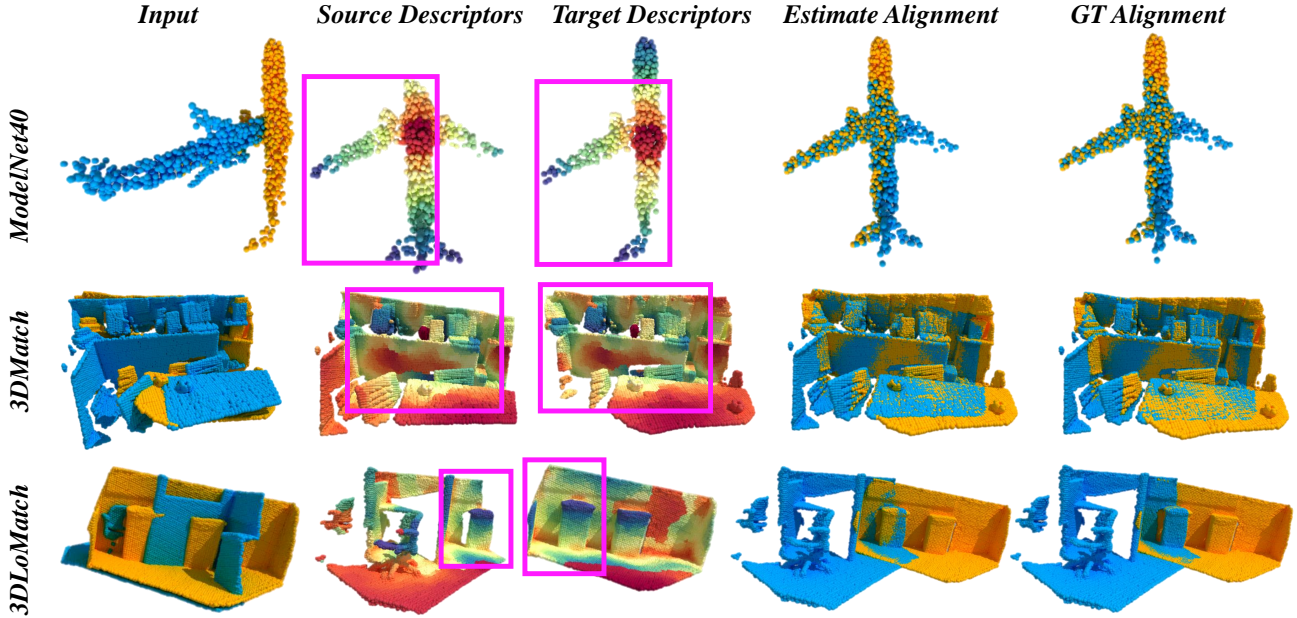
Fig. 7: **Qualitative Results.** We use t-SNE [5] to visualize the learned descriptors of source and target point clouds. In the rectangles, we roughly demonstrate the overlap regions.

TABLE 1: **Results on ModelNet40.** Best performance is highlighted in bold while the second best is marked with an underline. In "Unseen", 20 categories are used for training and the rest 20 for testing. In "Noise", all the categories are split into training and testing. Gaussian noise sampled from $\mathcal{N}(0, 0.01)$ and clipped to [-0.05, 0.05] is added to individual points in both training and testing. In "[0, 45°]", rotations along each axis are randomly sampled from [0, 45°] and translations are sampled from [-0.5, 0.5]. Rotations are enlarged to 180°in "[0, 180°]".

| | | Unseen | | | | | | Noise | | | | | |
| | | [0, 45°] | | | [0, 180°] | | | [0, 45°] | | | [0, 180°] | | |
Methods	#dim	RRE ↓	RTE ↓	RMSE ↓	RRE ↓	RTE ↓	RMSE ↓	RRE ↓	RTE ↓	RMSE ↓	RRE ↓	RTE ↓	RMSE ↓
PRNet [47]	1024	3.19°	0.028	0.036	91.94°	0.297	0.545	4.37°	0.034	0.045	95.80°	0.319	0.542
IDAM [49]	32	0.86°	0.005	0.007	16.17°	0.073	0.106	9.60°	0.052	0.084	71.06°	0.217	0.430
RPM [50]	1024	0.34°	0.004	0.004	8.78°	0.076	0.084	2.21°	0.013	0.018	23.58°	0.111	0.156
DCP [51]	1024	11.92°	0.076	0.119	67.39°	0.170	0.410	9.33°	0.070	0.097	73.61°	0.185	0.441
DeepGMR [52]	128	17.45°	0.074	0.130	49.23°	0.219	0.349	16.96°	0.068	0.120	68.68°	0.248	0.419
RPMNet [28]	96	0.60°	0.004	0.005	16.91°	0.079	0.127	3.52°	0.214	0.029	37.82°	0.132	0.250
GMCNet [29]	128	<u>0.026°</u>	<u>0.0002</u>	<u>0.0002</u>	**0.39°**	**0.002**	**0.003**	**0.94°**	<u>0.007</u>	**0.008**	18.13°	0.093	0.132
Predator [2]	96	<u>1.32°</u>	0.009	0.012	11.59°	<u>0.032</u>	0.058	3.33°	0.018	0.025	40.64°	0.110	0.207
CoFiNet [14]	32	2.30°	0.027	0.033	6.55°	0.033	<u>0.056</u>	3.06°	0.017	0.027	<u>14.33°</u>	<u>0.034</u>	<u>0.091</u>
RIGA	32	**0.004°**	**<0.0001**	**<0.0001**	<u>0.41°</u>	**0.002**	**0.003**	<u>1.15°</u>	**0.006**	<u>0.009</u>	**5.99°**	**0.008**	**0.029**

4.3 Real Scene Benchmarks: 3DMatch and 3DLoMatch

4.3.1 Datasets

3DMatch [43] collects 62 scenes, where 46 scenes are used for training, 8 for validation, and the rest 8 for testing. We use the processed data and split in [2], and evaluate RIGA on both 3DMatch [43] (>30% overlap) and 3DLoMatch [2] (10% ∼ 30% overlap) protocols. Additionally, we also follow [7], [3] to test on benchmarks with enlarged rotations to demonstrate the superiority of the inherent rotational invariance of our descriptors.

4.3.2 Metrics

We follow [2], [14] and use 3 metrics for evaluation: (1) *Inlier Ratio* (**IR**), which is the fraction of putative correspondences whose residual error is lower than a threshold $\tau_2 = 0.1m$ under the ground truth transformation, and (2) *Feature Matching Recall* (**FMR**) that counts the fraction of point cloud pairs whose **Inlier Ratio** is larger than a threshold $\tau_1 = 5\%$,

and (3) *Registration recall* (**RR**) that stands for the fraction of point cloud pairs whose **RMSE** between the estimated and ground truth transformation is smaller than a threshold $\tau_3 = 0.2m$. [1] Please refer to the Appendix for details.

4.3.3 Comparisons to the State-of-the-Art

In Tab. 3, we compare RIGA with 8 baseline methods. Specifically, 3DSN [8], SpinNet [3], and YOHO [4] are rotation-invariant approaches without global awareness. Predator [2], CoFiNet [14], and Lepard[2] [13] are globally-aware algorithms that are variant to rotations. We validate our method on both original and rotated benchmarks.[3] For

---

1. Instead of strictly following the criterion, we follow [2], [14] to calculate **RR** according to pre-computed correspondences defined on original 3DMatch/3DLoMatch.

2. We use the criterion in [2] and [14] to evaluate [13] and use all the correspondences without sampling following [13]

3. On rotated data, **RR** is calculated with **RMSE**<0.2m, which is different to **RR** on original data.
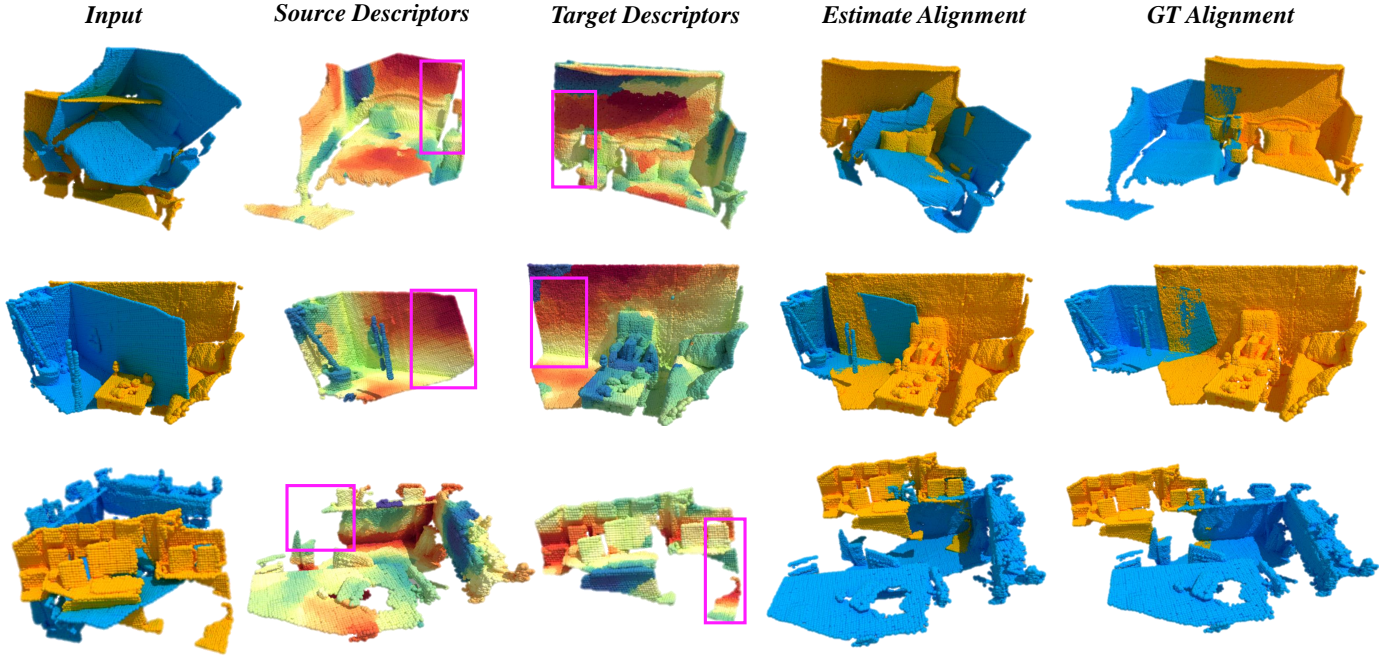
Fig. 8: **Failed cases on 3DLoMatch**. We use t-SNE [5] to visualize the learned descriptors of source and target point clouds. In the rectangles, we roughly demonstrate the overlap regions. The failed cases have reasonable descriptors but extremely limited overlap.

IR, RIGA significantly outperforms all the baselines on original 3DMatch and 3DLoMatch, which indicates RIGA learns more distinctive descriptors and extracts more reliable correspondences. When the benchmarks are further rotated, our superiority over others becomes more significant, which demonstrates the advantage of our rotational invariance by design. Notably, with larger rotations, only the performance of SpinNet [3], YOHO [4], and RIGA remains stable, which further proves the superiority of inherent rotational invariance over the learned one. For **FMR**, we perform the best on rotated data. When rotations are enlarged, especially on 3DLoMatch, the performance of all the methods except for RIGA and SpinNet [3] drops sharply. The performance drop of YOHO further demonstrates the aforementioned drawback of achieving rotational invariance via equivariance. Moreover, due to the lack of global awareness, SpinNet [3] falls behind Predator[2], CoFiNet [14], Lepard [13], and RIGA in terms of **FMR**, which supports the significance of being globally-aware. Finally, for **RR**, we perform on-par with CoFiNet [14] and Lepard [13] on original datasets, but again show our excellence when rotations are enlarged.

4.3.4 Detailed Results with Different Numbers of Samples
In Tab. 3, Tab. 4 and Fig. 9, we follow [2], [14] to show the performance with different numbers of sampled points/correspondences. The **IR** of CoFiNet [14] and RIGA increases when the number of samples decreases. This is because methods with the coarse-to-fine matching mechanism implicitly consider all the potential correspondences and sample the most confident ones for registration, while methods relying on uniform sub-sampling or keypoint detection only extract correspondences from sparsely-sampled nodes, whose repeatability is hard to guarantee especially with fewer samples. When the sample number is decreased

from 5,000 to 250, all the other metrics of CoFiNet and RIGA remain stable, while those of the others usually drop significantly, which further proves the excellence of the coarse-to-fine mechanism against fewer samples.

4.3.5 Scene-wise Results on 3DMatch and 3DLoMatch
We further detail the performance of RIGA with scene-wise results and 2 more metrics (**RRE** and **RTE**) in Tab. 5. The results further show the superiority of RIGA in scene-level registration.

4.4 Ablation Study
We ablate different parts of RIGA, including (1) *Local Description*, (2) *Global Description* and (3) *Attention Blocks* to assess the importance of each individual component. We use 3DMatch and 3DLoMatch, together with their rotated versions for ablation study. Detailed results are found in Tab. 6 for 3DLoMatch and Rotated 3DLoMatch, and in the Appendix for 3DMatch and Rotated 3DMatch.

4.4.1 Local Description
In the ablation of (1) *Local Description*, we replace our local PPF-based geometric description with two rotation-variant variants: (a) xyz - learning local descriptors from the raw 3D coordinates of all the points in the support area around each node; and (b) relative xyz - learning descriptors from relative 3D coordinates of points w.r.t. the central node of the support area. In both cases, the performance drops compared to the baseline RIGA, which indicates the power of our PPF signature-based geometric description. Moreover, we observe a more significant drop in performance in terms of **IR** and **FMR** when facing larger rotations which further demonstrates the importance of rotational invariance.

TABLE 2: **Comparisons to the State-of-the-Art on 3DMatch and 3DLoMatch.** Best performance is highlighted in bold while the second best is marked with an underline. In column "Rotated"[3], every point cloud pair is evaluated with # Samples=5,000[2] (in Tab. 3 and Tab. 4), and each point cloud is rotated individually with random rotations up to 360°along each axis. Our method significantly outperforms state-of-the-art methods on the rotated benchmarks.

# Samples	3DMatch		3DLoMatch	
	Origin	Rotated	Origin	Rotated
Inlier Ratio(%) ↑				
3DSN [8]	36.0	-	11.4	-
FCGF [9]	56.8	49.3	21.4	17.3
D3Feat [10]	39.0	37.7	13.2	12.1
SpinNet [3]	48.5	48.7	25.7	<u>25.7</u>
Predator [2]	58.0	52.8	26.7	22.4
YOHO [4]	<u>64.4</u>	<u>64.1</u>	25.9	23.2
CoFiNet [14]	49.8	46.8	24.4	21.5
Lepard [13]	58.6	53.7	<u>28.4</u>	<u>24.4</u>
RIGA	**68.4**	**68.5**	**32.1**	**32.1**
Feature Matching Recall(%) ↑				
3DSN [8]	95.0	-	63.6	-
FCGF [9]	97.4	96.9	76.6	73.3
D3Feat [10]	95.6	94.7	67.3	63.9
SpinNet [3]	97.4	97.4	75.5	75.2
Predator [2]	96.6	96.2	78.6	73.7
YOHO [4]	**98.2**	<u>97.8</u>	79.4	77.8
CoFiNet [14]	<u>98.1</u>	97.4	<u>83.1</u>	78.6
Lepard [13]	98.0	97.4	<u>83.1</u>	<u>79.5</u>
RIGA	97.9	**98.2**	**85.1**	**84.5**
Registration Recall(%) ↑				
3DSN [8]	78.4	-	33.0	-
FCGF [9]	85.1	90.3	40.1	58.6
D3Feat [10]	81.6	91.3	37.2	55.3
SpinNet [3]	88.8	**93.2**	58.2	61.8
Predator [2]	89.0	92.0	59.8	58.6
YOHO [4]	<u>90.8</u>	92.5	65.2	<u>66.8</u>
CoFiNet [14]	89.3	92.0	**67.5**	62.5
Lepard [13]	**92.7**	84.9	<u>65.4</u>	49.0
RIGA	89.3	<u>93.0</u>	65.1	**66.9**

TABLE 3: **Quantitative Results on 3DMatch and 3DLo-Match with Different Numbers of Samples.** Best performance is highlighted in bold while the second best is marked with an underline. # Samples is the number of sampled points or correspondences, following [2] and [14], respectively.

# Samples	3DMatch					3DLoMatch				
	5000	2500	1000	500	250	5000	2500	1000	500	250
Inlier Ratio(%) ↑										
3DSN [8]	36.0	32.5	26.4	21.5	16.4	11.4	10.1	8.0	6.4	4.8
FCGF [9]	56.8	54.1	48.7	42.5	34.1	21.4	20.0	17.2	14.8	11.6
D3Feat [10]	39.0	38.8	40.4	41.5	41.8	13.2	13.1	14.0	14.6	15.0
SpinNet [3]	48.5	46.2	40.8	35.1	29.0	25.7	23.7	20.6	18.2	13.1
Predator [2]	58.0	58.4	<u>57.1</u>	<u>54.1</u>	49.3	<u>26.7</u>	<u>28.1</u>	<u>28.3</u>	<u>27.5</u>	25.8
YOHO [4]	<u>64.4</u>	<u>60.7</u>	55.7	46.4	41.2	25.9	23.3	22.6	18.2	15.0
CoFiNet [14]	49.8	51.2	51.9	52.2	<u>52.2</u>	24.4	25.9	26.7	26.8	<u>26.9</u>
RIGA	**68.4**	**69.7**	**70.6**	**70.9**	**71.0**	**32.1**	**33.4**	**34.3**	**34.5**	**34.6**
Feature Matching Recall(%) ↑										
3DSN [8]	95.0	94.3	92.9	90.1	82.9	63.6	61.7	53.6	45.2	34.2
FCGF [9]	97.4	97.3	97.0	96.7	96.6	76.6	75.4	74.2	71.7	67.3
D3Feat [10]	95.6	95.4	94.5	94.1	93.1	67.3	66.7	67.0	66.7	66.5
SpinNet [3]	97.4	97.0	96.4	96.7	94.8	75.5	75.1	74.2	69.0	62.7
Predator [2]	96.6	96.6	96.5	96.3	96.5	78.6	77.4	76.3	75.7	75.3
YOHO [4]	**98.2**	97.6	97.5	<u>97.7</u>	96.0	79.4	78.1	76.3	73.8	69.1
CoFiNet [14]	<u>98.1</u>	**98.3**	**98.1**	**98.2**	**98.3**	<u>83.1</u>	<u>83.5</u>	<u>83.3</u>	<u>83.1</u>	<u>82.6</u>
RIGA	97.9	<u>97.8</u>	<u>97.7</u>	<u>97.7</u>	<u>97.6</u>	**85.1**	**85.0**	**85.1**	**84.3**	**85.1**
Registration Recall(%) ↑										
3DSN [8]	78.4	76.2	71.4	67.6	50.8	33.0	29.0	23.3	17.0	11.0
FCGF [9]	85.1	84.7	83.3	81.6	71.4	40.1	41.7	38.2	35.4	26.8
D3Feat [10]	81.6	84.5	83.4	82.4	77.9	37.2	42.7	46.9	43.8	39.1
SpinNet [3]	88.8	88.0	84.5	79.0	69.2	58.2	56.7	49.8	41.0	26.7
Predator [2]	89.0	<u>89.9</u>	**90.6**	88.5	86.6	59.8	61.2	62.4	60.8	58.1
YOHO [4]	**90.8**	**90.3**	<u>89.1</u>	<u>88.6</u>	84.5	<u>65.2</u>	<u>65.5</u>	63.2	56.5	48.0
CoFiNet [14]	<u>89.3</u>	88.9	88.4	87.4	<u>87.0</u>	**67.5**	**66.2**	<u>64.2</u>	<u>63.1</u>	<u>61.0</u>
RIGA	<u>89.3</u>	88.4	<u>89.1</u>	**89.0**	**87.7**	65.1	64.7	**64.5**	**64.1**	**61.8**

Similarly to [6], [28], we also concatenate PPF signatures with coordinates of points for local description in (c) and (d). This results in a better performance than the variants with only 3D coordinates but still perform slightly worse than the baseline RIGA. Thanks to the global awareness in RIGA, it is unnecessary to supplement PPF with global coordinates, as in (c), to incorporate global contexts. Pure local geometry which is rotation-invariant already promises good performance.

4.4.2 Global Description

We first ablate (2) *Global Description* by removing structural descriptors learned from our proposed global PPF signatures. As shown in (a), this significantly damages the performance especially in terms of **IR**, which proves the importance of informing local descriptors with global structural cues. To further prove the significance of our rotation-invariant structural description, we replace the structural descriptors in baseline RIGA with (b)xyz - learning global positional descriptors from the raw 3D coordinates of each node, and (c) relative xyz - learning global positional descriptors from the relative position of each node w.r.t.

TABLE 4: **Quantitative results on Rotated 3DMatch and 3DLoMatch with Different Numbers of Samples.** Best performance is highlighted in bold while the second best is marked with an underline. Each point cloud is rotated individually with random rotations up to 360°along each axis.

# Samples	3DMatch					3DLoMatch				
	5000	2500	1000	500	250	5000	2500	1000	500	250
Inlier Ratio(%) ↑										
FCGF [9]	49.3	47.1	42.5	37.4	30.6	17.3	16.4	14.6	12.5	10.2
D3Feat [10]	37.7	37.7	37.0	36.0	34.6	12.1	12.1	11.9	11.7	11.2
SpinNet [3]	48.7	46.0	40.6	35.1	29.0	<u>25.7</u>	<u>23.9</u>	20.8	17.9	15.6
Predator [2]	52.8	53.4	52.5	<u>50.0</u>	45.6	22.4	23.5	23.0	23.2	21.6
YOHO [4]	<u>64.1</u>	<u>60.4</u>	<u>53.5</u>	46.3	36.9	23.2	23.2	19.2	15.7	12.1
CoFiNet [14]	46.8	48.2	49.0	49.3	<u>49.3</u>	21.5	22.8	<u>23.6</u>	<u>23.8</u>	<u>23.8</u>
RIGA	**68.5**	**69.8**	**70.7**	**71.0**	**71.2**	**32.1**	**33.5**	**34.3**	**34.7**	**35.0**
Feature Matching Recall(%) ↑										
FCGF [9]	96.9	96.9	96.2	95.9	94.5	73.3	73.4	71.0	68.8	64.5
D3Feat [10]	94.7	95.1	94.3	93.8	92.3	63.9	64.6	63.0	62.1	59.6
SpinNet [3]	97.4	97.4	96.7	96.5	94.1	75.2	74.9	72.6	69.2	61.8
Predator [2]	96.2	96.2	96.6	96.0	96.0	73.7	74.2	75.0	74.8	73.5
YOHO [4]	<u>97.8</u>	<u>97.8</u>	<u>97.4</u>	<u>97.6</u>	96.4	77.8	77.8	76.3	73.9	67.3
CoFiNet [14]	97.4	97.4	97.2	97.2	<u>97.3</u>	<u>78.6</u>	<u>78.8</u>	<u>79.2</u>	<u>78.9</u>	<u>79.2</u>
RIGA	**98.2**	**98.2**	**98.2**	**98.0**	**98.1**	**84.5**	**84.6**	**84.5**	**84.2**	**84.4**
Registration Recall(%) ↑										
FCGF [9]	90.3	91.2	90.4	87.8	83.3	58.6	58.7	54.7	44.8	34.7
D3Feat [10]	91.3	90.3	88.4	85.2	80.8	55.3	53.5	47.9	43.6	33.5
SpinNet [3]	**93.2**	**93.2**	91.1	87.4	77.0	61.8	59.1	53.1	44.1	30.7
Predator [2]	92.0	92.8	92.0	**92.2**	89.5	58.6	59.5	60.4	58.6	55.8
YOHO [4]	92.5	92.3	<u>92.4</u>	90.2	87.4	<u>66.8</u>	<u>67.1</u>	<u>64.5</u>	58.2	44.8
CoFiNet [14]	92.0	91.4	91.0	90.3	<u>89.6</u>	62.5	60.9	60.9	<u>59.9</u>	<u>56.5</u>
RIGA	<u>93.0</u>	<u>93.0</u>	**92.6**	<u>91.8</u>	**92.3**	**66.9**	**67.6**	**67.0**	**66.5**	**66.2**

the other nodes in the same frame. Moreover, we also follow [23] to learn descriptors from node coordinates projected by sinusoidal functions [5] in (d). The decreased performance of all the variants further confirms the superiority of our design of encoding structural descriptors from global PPF signatures.

TABLE 5: **Scene-Wise Results on 3DMatch and 3DLoMatch with #Samples=5,000.** Best performance is highlighted in bold while the second best is marked with an underline.

Method	3DMatch									3DLoMatch								
	Kitchen	Home_1	Home_2	Hotel_1	Hotel_2	Hotel_3	Study	Lab	Mean	Kitchen	Home_1	Home_2	Hotel_1	Hotel_2	Hotel_3	Study	Lab	Mean
Registration Recall(%)↑																		
3DSN [8]	90.6	90.6	65.4	89.6	82.1	80.8	68.4	60.0	78.4	51.4	25.9	44.1	41.1	30.7	36.6	14.0	20.3	33.0
FCGF [9]	98.0	94.3	68.6	96.7	91.0	84.6	76.1	71.1	85.1	60.8	42.2	53.6	53.1	38.0	26.8	16.1	30.4	40.1
D3Feat [10]	96.0	86.8	67.3	90.7	88.5	80.8	78.2	64.4	81.6	49.7	37.2	47.3	47.8	36.5	31.7	15.7	31.9	59.8
Predator [2]	97.6	97.2	74.8	98.9	96.2	88.5	85.9	73.3	89.0	71.5	58.2	60.8	77.5	64.2	61.0	45.8	39.1	59.8
CoFiNet [14]	96.4	99.1	73.6	95.6	91.0	84.6	89.7	84.4	89.3	76.7	66.7	64.0	81.3	65.0	63.4	53.4	69.6	67.5
RIGA	97.8	93.4	76.7	98.4	93.6	84.6	85.9	84.4	89.3	77.8	60.6	63.5	79.4	62.0	63.4	48.7	65.2	65.1
Relative Rotation Error(°)↓																		
3DSN [8]	1.926	1.843	2.324	2.041	1.952	2.908	2.296	2.301	2.199	3.020	3.898	3.427	3.196	3.217	3.328	4.325	3.814	3.528
FCGF [9]	1.767	1.849	2.210	1.867	1.667	2.417	2.024	1.792	1.949	2.904	3.229	3.277	2.768	2.801	2.822	3.372	4.006	3.147
D3Feat [10]	2.016	2.029	2.425	1.990	1.967	2.400	2.346	2.115	2.161	3.226	3.492	3.373	3.330	3.165	2.972	3.708	3.619	3.361
Predator [2]	1.861	1.806	2.473	2.045	1.600	2.458	2.067	1.926	2.029	3.079	2.637	3.220	2.694	2.907	3.390	3.046	3.412	3.048
CoFiNet [14]	1.910	1.835	2.316	1.767	1.753	1.639	2.527	2.345	2.011	3.213	3.119	3.711	2.842	2.897	3.194	4.126	3.138	3.280
RIGA	1.789	1.538	1.981	1.677	1.598	1.935	1.833	2.033	1.798	2.987	2.722	3.313	2.743	2.956	2.439	3.836	3.135	3.016
Relative Translation Error(m)↓																		
3DSN [8]	0.059	0.070	0.079	0.065	0.074	0.062	0.093	0.065	0.071	0.082	0.098	0.096	0.101	0.080	0.089	0.158	0.120	0.103
FCGF [9]	0.053	0.056	0.071	0.062	0.061	0.055	0.082	0.090	0.066	0.084	0.097	0.076	0.101	0.084	0.077	0.144	0.140	0.100
D3Feat [10]	0.053	0.065	0.080	0.064	0.078	0.049	0.083	0.064	0.067	0.088	0.101	0.086	0.099	0.092	0.075	0.146	0.135	0.103
Predator [2]	0.048	0.055	0.070	0.073	0.060	0.065	0.080	0.063	0.064	0.081	0.080	0.084	0.099	0.096	0.077	0.101	0.130	0.093
CoFiNet [14]	0.047	0.059	0.063	0.063	0.058	0.044	0.087	0.075	0.062	0.080	0.078	0.078	0.099	0.086	0.077	0.131	0.123	0.094
RIGA	0.044	0.048	0.056	0.060	0.059	0.040	0.071	0.071	0.056	0.078	0.082	0.085	0.094	0.082	0.059	0.116	0.114	0.089

TABLE 6: **Ablation Study on 3DLoMatch and Rotated 3DLoMatch with # Samples = 5,000.** In the brackets are the changes compared to baseline RIGA.

| Ablation Part	Models	3DLoMatch			3DLoMatch (Rotated)		
		IR(%) ↑	FMR(%) ↑	RR(%) ↑	IR(%) ↑	FMR(%) ↑	RR(%)↑
(0) *None*	RIGA (*Baseline*)	32.1	85.1	65.1	32.1	84.5	66.9
(1) *Local Description*	(a) xyz	20.8(−11.3)	77.5(−7.60)	56.0(−9.10)	20.2(−11.9)	76.2(−8.30)	57.4(−9.50)
	(b) relative xyz	25.7(−6.40)	79.6(−5.50)	58.5(−6.60)	24.9(−7.20)	79.9(−4.60)	59.9(−7.00)
	(c) xyz + PPF	31.1(−1.00)	85.1(+0.00)	65.3(+0.20)	31.1(−1.00)	83.6(−0.90)	66.5(−0.40)
	(d) relative xyz + PPF	30.7(−1.40)	83.6(−1.50)	62.5(−2.60)	30.6(−1.50)	83.1(−1.40)	64.5(−2.40)
(2) *Global Description*	(a) none	13.8(−18.3)	75.4(−9.70)	61.1(−4.00)	13.9(−18.2)	76.0(−8.50)	66.0(−0.90)
	(b) xyz	18.6(−13.5)	81.3(−3.80)	65.1(+0.00)	18.5(−13.6)	80.5(−4.00)	65.8(−1.10)
	(c) relative xyz	15.1(−17.0)	77.5(−7.60)	62.8(−2.30)	14.8(−17.3)	75.7(−8.80)	65.2(−1.70)
	(d) xyz+sinusoidal [23]	15.1(−17.0)	76.7(−8.40)	64.6(−0.50)	15.1(−17.0)	78.3(−6.20)	66.5(−0.40)
(3) *Attention Blocks*	(a) $K=0$	10.2(−21.9)	60.8(−24.3)	50.0(−15.1)	10.3(−21.8)	60.2(−24.3)	53.6(−13.3)
	(b) $K=1$	16.6(−15.5)	77.1(−8.00)	63.0(−2.10)	16.6(−15.5)	77.8(−6.70)	66.1(−0.80)
	(c) $K=3$	24.9(−7.20)	82.2(−2.90)	65.1(+0.00)	25.0(−7.10)	82.4(−2.10)	66.6(−0.30)
	(d) $K=10$	32.5(+0.40)	83.6(−1.50)	63.6(−1.50)	32.4(+0.30)	83.4(−1.10)	66.8(−0.10)
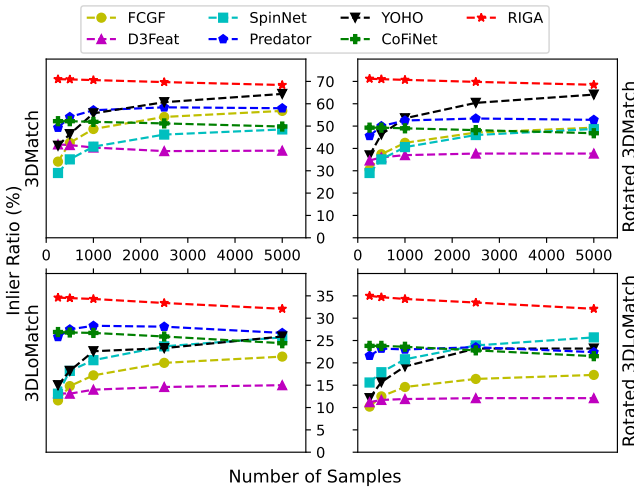
Fig. 9: **Inlier Ratio (IR) with Different Numbers of Samples.** RIGA achieves the best performance on all the datasets. Notably, the performance of RIGA increases when the number of sampled correspondences decreases, which further demonstrates the superiority of our coarse-to-fine mechanism for correspondence extraction.

4.4.3 Attention Blocks

To emphasize the importance of global awareness, we ablate RIGA with different number of (3) *Attention Blocks*. In (a), we remove all the attention blocks ($K=0$) and only use the globally-informed descriptors, which leads to a sharp decrease of the performance. This proves the significance of global awareness obtained from learned global contexts. When we increase the number of attention blocks to (b) $K=1$ and (c) $K=3$, the performance increases correspondingly, though it does not reach the baseline performance with $K=6$. This observation indicates that stronger global awareness improves the overall performance. However, when we keep including more and more Attention Blocks in (d) $K=10$, the performance only stays on-par with RIGA baseline, indicating that using 6 Attention Blocks is a proper option with good performance.

5 CONCLUSION

In this paper, we introduce RIGA with a ViT architecture that learns both rotation-invariant and globally-aware descriptors, upon which correspondences are established in a coarse-to-fine manner for point cloud registration. We learn from rotation-invariant PPFs for encoding local geometry and further introduce global PPF signatures to encode a node-specific structural description of the whole scene. The

structural descriptors learned from global PPF signatures strengthen local descriptors with the global 3D structures in a rotation-invariant fashion. The distinctiveness of descriptors is further enhanced in the consecutive attention blocks with the learned geometric context across the whole scene. The coarse-to-fine mechanism is further leveraged to establish reliable correspondences upon our powerful RIGA descriptors. Experimental results confirm the effectiveness of our approach on both object and scene-level data. We hope our work can inspire more research looking toward the joint rotational invariance and distinctiveness of descriptors in point cloud registration.

6 APPENDIX

In this Appendix, we first detail related metrics in Sec. 6.1. We then demonstrate the inference speed of RIGA in Sec. 6.2 and detail the ablation study on 3DMatch [43] and Rotated 3DMatch in Sec. 6.3. We further use KITTI [44] to demonstrate our robustness against poor normal estimation in Sec. 6.4. Finally, more quantitative results are illustrated in Sec. 6.5.

6.1 Detailed Metrics

Relative Rotation and Translation Errors. Given the estimated rotation $\overline{\mathbf{R}} \in SO(3)$ and translation $\overline{\mathbf{t}} \in \mathbb{R}^3$ between a pair of point clouds $(\mathcal{X}, \mathcal{Y})$, the *Relative Rotation Error* and the *Relative Translation Error* w.r.t. the ground truth rotation $\mathbf{R} \in SO(3)$ and translation $\mathbf{t} \in \mathbb{R}^3$ are computed as:

$$\mathbf{RRE}(\mathcal{X}, \mathcal{Y}) = arccos(\frac{trace(\mathbf{R}^T \overline{\mathbf{R}}) - 1}{2}), and \tag{24}$$
$$\mathbf{RTE}(\mathcal{X}, \mathcal{Y}) = \|\mathbf{t} - \overline{\mathbf{t}}\|_2,$$

respectively.

Root-Mean-Square Error. Given the estimated transformation $\overline{\mathbf{T}} \in SE(3)$ and the ground truth transformation $\mathbf{T}$ between a pair of point clouds $(\mathcal{X}, \mathcal{Y})$, there can be two ways to calculate *Root-Mean-Square Error* (**RMSE**). According to [52], the first way to calculate **RMSE** reads as:

$$\mathbf{RMSE_1}(\mathcal{X}, \mathcal{Y}) = \frac{1}{|\mathcal{X}|}\sqrt{\sum_{\mathbf{x} \in \mathcal{X}}\|\overline{\mathbf{T}}(\mathbf{x}) - \mathbf{T}(\mathbf{x})\|_2^2}, \tag{25}$$

which is used for the experiments on ModelNet [42] and the calculation of *Registration Recall* on rotated 3DMatch [43] and rotated 3DLoMatch [2]. Additionally, we follow [2] to calculate $\mathbf{RMSE_2}$, upon which the *Registration Recall* on 3DMatch and 3DLoMatch is further defined. $\mathbf{RMSE_2}$ is calculated as:

$$\mathbf{RMSE_2}(\mathcal{X}, \mathcal{Y}) = \sqrt{\frac{1}{|\mathcal{C}^*|}\sum_{(\mathbf{x},\mathbf{y}) \in \mathcal{C}^*}\|\overline{\mathbf{T}}(\mathbf{x}) - \mathbf{y}\|_2^2}, \tag{26}$$

where $\mathcal{C}^*$ is a ground truth correspondence set.

Inlier Ratio. *Inlier Ration* (**IR**) measures the fraction of putative correspondences $(\mathbf{x}, \mathbf{y}) \in \hat{\mathcal{C}}$ s.t. $\|\mathbf{T}(\mathbf{x}) - \mathbf{y}\|_2$ is within a threshold $\tau_1 = 10cm$, where $\mathbf{T}$ stands for the ground truth transformation between point clouds $\mathcal{X}$ and $\mathcal{Y}$.

TABLE 7: **Runtime.** All the reported time is averaged over the whole 3DMatch testing set, which consists of 1,623 point cloud pairs. "Desc" reports the runtime for description, i.e., from data loading to the generation of descriptors. "Reg" reports the time for registration, i.e., from the generated descriptors to the estimation of rigid transformation via RANSAC [1]. These two parts of time sum to "Total".

Method	Desc (s)↓	Reg (s)↓	Total (s)↓
SpinNet [3]	44.92	-	>44.92
Predator [2]	0.506	0.677	1.183
CoFiNet [14]	**0.145**	**0.043**	**0.188**
RIGA (*Ours*)	0.731	0.101	0.832

The **IR** of a single point cloud pair $(\mathcal{X}, \mathcal{Y})$ with a putative correspondence set $\hat{\mathcal{C}}$ is defined as:

$$\mathbf{IR}(\mathcal{X}, \mathcal{Y}) = \frac{1}{|\hat{\mathcal{C}}|}\sum_{(\mathbf{x},\mathbf{y}) \in \hat{\mathcal{C}}}\mathbb{1}(\|\mathbf{T}(\mathbf{x}) - \mathbf{y}\|_2 < \tau_1), \tag{27}$$

where $\mathbb{1}(\cdot)$ denotes the indicator function.

Feature Matching Recall. *Feature Matching Recall* (**FMR**) counts the fraction of point cloud pairs $(\mathcal{X}, \mathcal{Y})$ that satisfies $\mathbf{IR}(\mathcal{X}, \mathcal{Y}) > \tau_2$, which is set to 5% in our experiments. Given a dataset $\mathcal{H}$ consisting of $|\mathcal{H}|$ point cloud pairs, the **FMR** is computed as:

$$\mathbf{FMR}(\mathcal{H}) = \frac{1}{|\mathcal{H}|}\sum_{(\mathcal{X},\mathcal{Y}) \in \mathcal{H}}\mathbb{1}(\mathbf{IR}(\mathcal{X}, \mathcal{Y}) > \tau_2). \tag{28}$$

Registration Recall. *Registration Recall* (**RR**) that measures the fraction of successfully registered point cloud pairs directly evaluates the performance of a method on the task of point cloud registration. More specifically, it counts the fraction of point cloud pairs $(\mathcal{X}, \mathcal{Y})$ that satisfies $\mathbf{RMSE_2}(\mathcal{X}, \mathcal{Y}) < \tau_3$, where τ_3 is set to 0.2m in our experiments. Given a dataset with $|\mathcal{H}|$ point cloud pairs, **RR** is defined as:

$$\mathbf{RR}(\mathcal{H}) = \frac{1}{|\mathcal{H}|}\sum_{(\mathcal{X},\mathcal{Y}) \in \mathcal{H}}\mathbb{1}(\mathbf{RMSE_2}(\mathcal{X}, \mathcal{Y}) < \tau_3). \tag{29}$$

6.2 Runtime Analysis

We test all the following approaches on a machine with "AMD Ryzen 7 5800X @ 3.80GHZ × 8" CPU and "NVIDIA GeForce RTX 3090" GPU. In Tab. 7 we compare RIGA with 3 state-of-the-art methods in terms of runtime. Among all the baselines, SpinNet [3] is a patch-based rotation-invariant method, while Predator [2] and CoFiNet [14] are globally-aware models with fully-convolutional encoder-decoder architectures. As RIGA use a ViT architecture that starts from the description of local regions, when compared to Predator and CoFiNet, it takes more time to generate descriptors. However, RIGA generates descriptors much faster than SpinNet, as our global awareness simplifies the feature engineering on local regions, and our mechanism in tackling the repeatability issues significantly reduces the number of required local regions. For registration time, as we adopt a coarse-to-fine strategy, the runtime is significantly reduced when compared to Predator. Moreover, we use the second least total time among all the methods, which demonstrates our efficiency for the task of point cloud registration.

Car

Remote Noisy Objects

Tree

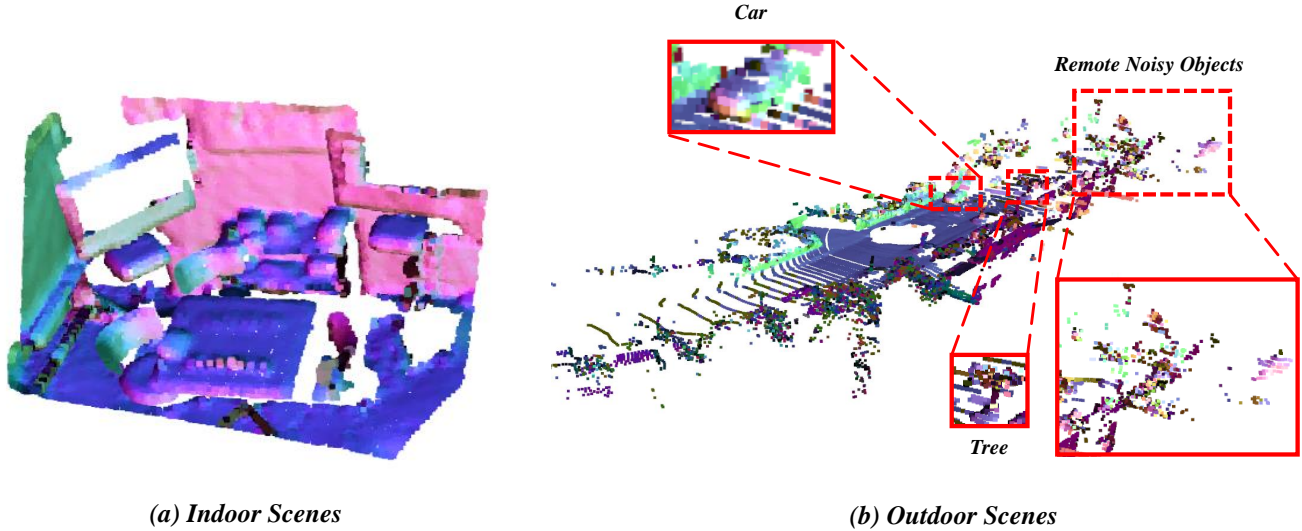*(a) Indoor Scenes*   *(b) Outdoor Scenes*

Fig. 10: **Demonstration of the Quality of Normal Estimation.** Normals are estimated by using Open3D [48] and are color-coded for visualization. The indoor scene in column (a) is from 3DMatch [43], while the outdoor scene in column (b) is from KITTI [44]. For indoor scenes, the normal estimation is accurate, i.e., the colors are smooth in the visualization. However, the quality of estimated normals in outdoor scenes is much worse. Although the estimated normals are not bad for the "Car" which is represented clearly by points with less noise, the normals of the "Tree" are worse, due to its complex geometry and noisy representation. Moreover, the objects that are far away from the LiDAR and roughly represented by sparse points are hard to recognize and with the worst normal quality.

TABLE 8: **Ablation study on 3DMatch and Rotated 3DMatch.** In the brackets are the changes compared to baseline RIGA. # Samples = 5,000.

Ablation Part	Models	3DMatch			3DMatch (Rotated)		
		IR(%) ↑	FMR(%) ↑	RR(%) ↑	IR(%) ↑	FMR(%) ↑	RR(%) ↑
(0) *None*	RIGA (*Baseline*)	68.4	97.9	89.3	68.5	98.2	93.0
(1) *Local Description*	(a) xyz	53.7(−14.7)	96.1(−1.80)	86.8(−2.50)	52.7(−15.8)	95.8(−2.40)	89.1(−3.10)
	(b) relative xyz	60.9(−7.50)	97.2(−0.70)	87.5(−1.80)	60.0(−8.50)	96.4(−1.80)	90.3(−2.70)
	(c) xyz + PPF	66.3(−2.10)	98.2(+0.30)	88.5(−0.80)	65.9(−2.60)	98.1(−0.10)	92.4(−0.60)
	(d) relative xyz + PPF	66.8(−1.60)	97.5(−0.40)	87.7(−1.60)	66.7(−1.80)	97.4(−0.80)	92.1(−0.90)
(2) *Global Description*	(a) none	34.9(−33.5)	97.0(−0.90)	88.1(−1.20)	35.0(−33.5)	97.0(−1.20)	92.8(−0.20)
	(b) xyz	42.3(−26.1)	97.8(−0.10)	87.7(−1.60)	42.3(−26.2)	97.6(−0.60)	92.3(−0.70)
	(c) relative xyz	37.2(−31.2)	97.0(−0.90)	88.0(−1.30)	37.0(−31.5)	96.8(−1.40)	93.3(+0.30)
	(d) xyz+sinusoidal [23]	37.1(−31.3)	97.1(−0.80)	89.8(+0.50)	37.1(−31.4)	97.6(−0.60)	93.0(+0.00)
(3) *Attention Blocks*	(a) K=0	29.7(−38.7)	94.2(−3.70)	83.0(−6.30)	29.5(−39.0)	94.2(−4.00)	89.3(−3.70)
	(b) K=1	43.7(−24.7)	97.4(−0.50)	90.1(+0.80)	43.7(−24.8)	97.3(−0.90)	93.4(+0.40)
	(c) K=3	58.1(−10.3)	97.7(−0.20)	88.8(−0.50)	58.4(−10.1)	98.1(−0.10)	92.7(−0.30)
	(d) K=10	68.5(+0.10)	98.1(+0.20)	89.0(−0.30)	68.4(−0.10)	97.9(−0.30)	92.2(−0.80)

TABLE 9: **Quantitative comparisons on KITTI.** Best performance is highlighted in bold.

Method	RTE(cm)↓	RRE(°)↓	RR(%)↑
3DFeat-Net [53]	25.9	0.57	96.0
FCGF [9]	9.5	0.30	96.6
D3Feat [10]	7.2	0.30	**99.8**
SpinNet [3]	9.9	0.47	99.1
Predator [2]	**6.8**	**0.27**	**99.8**
CoFiNet [14]	8.5	0.41	**99.8**
RIGA (*Ours*)	13.5	0.45	99.1

6.3 Ablation Studies on 3DMatch/Rotated 3DMatch

In Tab. 8, we show the ablation study on 3DMatch and rotated 3DMatch with the same setting as in the ablation study of 3DLoMatch and rotated 3DLoMatch in the main paper. Different variants behave similarly , which further illustrates the significance of each individual part of RIGA.

6.4 Robustness against Poor Normal Estimation

As our inherent rotational invariance is affected by the quality of the estimated normals, we further conduct extensive experiments on KITTI [44] which consists of outdoor scans from LiDAR to prove the robustness of our RIGA descriptors against poor normal estimation. The estimated normals of both indoor and outdoor scenarios are visualized in Fig. 10 to show the poor normal estimation for outdoor

scenes compared to indoor ones. Under this circumstance, as shown in Tab. 9, although RIGA is affected by the poor normal quality, it still performs on par with those state-of-the-art methods in terms of three different metrics.

6.5 More Qualitative Results.

More qualitative results on both ModetNet40 and 3DMatch/3DLoMatch can be found in Fig. 11 and Fig. 12, respectively. In each figure, the first column gives a pair of unaligned point clouds, where the source point cloud is presented as blue and the target point cloud is shown in yellow. The second and third columns illustrate the RIGA descriptors visualized by t-SNE [5] for source and target point clouds, respectively. The forth column demonstrates the estimated alignment, while the last column provides the ground truth one.

REFERENCES

[1] M. A. Fischler and R. C. Bolles, "Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography," *Communications of the ACM*, 1981. 1, 8, 13

[2] S. Huang, Z. Gojcic, M. Usvyatsov, A. Wieser, and K. Schindler, "Predator: Registration of 3d point clouds with low overlap," in *CVPR*, 2021. 1, 2, 3, 4, 5, 7, 8, 9, 10, 11, 12, 13, 14

[3] S. Ao, Q. Hu, B. Yang, A. Markham, and Y. Guo, "Spinnet: Learning a general surface descriptor for 3d point cloud registration," in *CVPR*, 2021. 1, 2, 3, 4, 5, 9, 10, 11, 13, 14

[4] H. Wang, Y. Liu, Z. Dong, W. Wang, and B. Yang, "You only hypothesize once: Point cloud registration with rotation-equivariant descriptors," *arXiv preprint arXiv:2109.00182*, 2021. 1, 2, 3, 5, 9, 10, 11

[5] L. Van der Maaten and G. Hinton, "Visualizing data using t-sne." *Journal of Machine Learning Research*, 2008. 2, 9, 10, 11, 15, 16, 17

[6] H. Deng, T. Birdal, and S. Ilic, "Ppfnet: Global context aware local features for robust 3d point matching," in *CVPR*, 2018. 1, 2, 4, 5, 11

[7] ——, "Ppf-foldnet: Unsupervised learning of rotation invariant 3d local descriptors," in *ECCV*, 2018. 1, 2, 3, 9

[8] Z. Gojcic, C. Zhou, J. D. Wegner, and A. Wieser, "The perfect match: 3d point cloud matching with smoothed densities," in *CVPR*, 2019. 1, 2, 3, 5, 9, 11, 12

[9] C. Choy, J. Park, and V. Koltun, "Fully convolutional geometric features," in *ICCV*, 2019. 1, 2, 4, 11, 12, 14

[10] X. Bai, Z. Luo, L. Zhou, H. Fu, L. Quan, and C.-L. Tai, "D3feat: Joint learning of dense detection and description of 3d local features," in *CVPR*, 2020. 1, 2, 4, 7, 11, 12, 14

[11] M. Saleh, S. Dehghani, B. Busam, N. Navab, and F. Tombari, "Graphite: Graph-induced feature extraction for point cloud registration," in *3DV*, 2020. 1, 2, 3

[12] M. Saleh, S.-C. Wu, L. Cosmo, N. Navab, B. Busam, and F. Tombari, "Bending graphs: Hierarchical shape matching using gated optimal transport," *arXiv preprint arXiv:2202.01537*, 2022. 1, 3

[13] Y. Li and T. Harada, "Lepard: Learning partial point cloud matching in rigid and deformable scenes," in *CVPR*, 2022. 1, 2, 3, 4, 5, 9, 10, 11

[14] H. Yu, F. Li, M. Saleh, B. Busam, and S. Ilic, "Cofinet: Reliable coarse-to-fine correspondences for robust pointcloud registration," in *NeurIPS*, 2021. 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14

[15] C. R. Qi, H. Su, K. Mo, and L. J. Guibas, "Pointnet: Deep learning on point sets for 3d classification and segmentation," in *CVPR*, 2017. 2, 4, 5, 8

[16] C. R. Qi, L. Yi, H. Su, and L. J. Guibas, "Pointnet++: Deep hierarchical feature learning on point sets in a metric space," in *NeurIPS*, 2017. 2, 4

[17] H. Thomas, C. R. Qi, J.-E. Deschaud, B. Marcotegui, F. Goulette, and L. J. Guibas, "Kpconv: Flexible and deformable convolution for point clouds," in *ICCV*, 2019. 2

[18] J. Tang, D. Xu, K. Jia, and L. Zhang, "Learning parallel dense correspondence from spatio-temporal descriptors for efficient and robust 4d reconstruction," in *CVPR*, 2021. 2

[19] J. Hou, B. Graham, M. Nießner, and S. Xie, "Exploring data-efficient 3d scene understanding with contrastive scene contexts," in *CVPR*, 2021. 2

[20] R. B. Rusu, N. Blodow, Z. C. Marton, and M. Beetz, "Aligning point cloud views using persistent feature histograms," in *IROS*, 2008. 2, 3

[21] R. B. Rusu, N. Blodow, and M. Beetz, "Fast point feature histograms (fpfh) for 3d registration," in *ICRA*, 2009. 2, 3

[22] B. Drost, M. Ulrich, N. Navab, and S. Ilic, "Model globally, match locally: Efficient and robust 3d object recognition," in *CVPR*, 2010. 2, 3, 4

[23] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," in *NeurIPS*, 2017. 2, 3, 5, 7, 11, 12, 14

[24] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly *et al.*, "An image is worth 16x16 words: Transformers for image recognition at scale," *arXiv preprint arXiv:2010.11929*, 2020. 2

[25] F. Tombari, S. Salti, and L. D. Stefano, "Unique signatures of histograms for local surface description," in *ECCV*, 2010. 3

[26] Y. Guo, F. Sohel, M. Bennamoun, M. Lu, and J. Wan, "Rotational projection statistics for 3d local surface description and object recognition," *IJCV*, 2013. 3

[27] Y. Yang, C. Feng, Y. Shen, and D. Tian, "Foldingnet: Point cloud auto-encoder via deep grid deformation," in *CVPR*, 2018. 3

[28] Z. J. Yew and G. H. Lee, "Rpm-net: Robust point matching using learned features," in *CVPR*, 2020. 3, 9, 11

[29] L. Pan, Z. Cai, and Z. Liu, "Robust partial-to-partial point cloud registration in a full range," *arXiv preprint arXiv:2111.15606*, 2021. 3, 8, 9

[30] T. Birdal and S. Ilic, "Point pair features based object detection and pose estimation revisited," in *3DV*, 2015. 3, 4

[31] ——, "Cad priors for accurate and flexible instance reconstruction," in *ICCV*, 2017. 3

[32] S. Hinterstoisser, V. Lepetit, N. Rajkumar, and K. Konolige, "Going further with point pair features," in *ECCV*, 2016. 3

[33] Y. Wang, Y. Sun, Z. Liu, S. E. Sarma, M. M. Bronstein, and J. M. Solomon, "Dynamic graph cnn for learning on point clouds," *ACM TOG*, 2019. 3

[34] Z. J. Yew and G. H. Lee, "Regtr: End-to-end point cloud correspondences with transformers," in *CVPR*, 2022. 3, 4

[35] Z. Qin, Z. Yu, C. Wang, Y. Guo, Y. Peng, and K. Xu, "Geometric transformer for fast and robust point cloud registration," *arXiv preprint arXiv:2202.06688*, 2022. 3, 4, 7

[36] H. Hoppe, T. DeRose, T. Duchamp, J. McDonald, and W. Stuetzle, "Surface reconstruction from unorganized points," in *SIGGRAPH*, 1992. 4

[37] P.-E. Sarlin, D. DeTone, T. Malisiewicz, and A. Rabinovich, "Superglue: Learning feature matching with graph neural networks," in *CVPR*, 2020. 5, 6, 7

[38] R. Sinkhorn and P. Knopp, "Concerning nonnegative matrices and doubly stochastic matrices," *Pacific Journal of Mathematics*, 1967. 6, 8

[39] J. Li and G. H. Lee, "Usip: Unsupervised stable interest point detection from 3d point clouds," in *ICCV*, 2019. 6

[40] Y. Sun, C. Cheng, Y. Zhang, C. Zhang, L. Zheng, Z. Wang, and Y. Wei, "Circle loss: A unified perspective of pair similarity optimization," in *CVPR*, 2020. 7

[41] F. Schroff, D. Kalenichenko, and J. Philbin, "Facenet: A unified embedding for face recognition and clustering," in *CVPR*, 2015. 7

[42] Z. Wu, S. Song, A. Khosla, F. Yu, L. Zhang, X. Tang, and J. Xiao, "3d shapenets: A deep representation for volumetric shapes," in *CVPR*, 2015. 8, 13

[43] A. Zeng, S. Song, M. Nießner, M. Fisher, J. Xiao, and T. Funkhouser, "3dmatch: Learning local geometric descriptors from rgb-d reconstructions," in *CVPR*, 2017. 8, 9, 13, 14

[44] A. Geiger, P. Lenz, and R. Urtasun, "Are we ready for autonomous driving? the kitti vision benchmark suite," in *CVPR*, 2012. 8, 13, 14

[45] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga *et al.*, "Pytorch: An imperative style, high-performance deep learning library," in *NeurIPS*, 2019. 8

Fig. 11: **More Qalitative Results on ModelNet40.** We use t-SNE [5] to visualize the learned descriptors of source and target point clouds.

[46] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014. 8

[47] Y. Wang and J. M. Solomon, "Prnet: Self-supervised learning for partial-to-partial registration," in *NeurIPS*, 2019. 8, 9

[48] Q.-Y. Zhou, J. Park, and V. Koltun, "Open3d: A modern library for 3d data processing," *arXiv preprint arXiv:1801.09847*, 2018. 8, 14

[49] J. Li, C. Zhang, Z. Xu, H. Zhou, and C. Zhang, "Iterative distance-aware similarity matrix convolution with mutual-supervised point elimination for efficient point cloud registration," in *ECCV*, 2020. 9

[50] K. Fu, S. Liu, X. Luo, and M. Wang, "Robust point cloud registration framework based on deep graph matching," in *CVPR*, 2021.

Fig. 12: **More qualitative results on 3DMatch/3DLoMatch.** We use t-SNE [5] to visualize the learned descriptors of source and target point clouds. In the rectangles, we roughly demonstrate the overlap regions.

9

[51] Y. Wang and J. M. Solomon, "Deep closest point: Learning representations for point cloud registration," in *ICCV*, 2019. 9

[52] W. Yuan, B. Eckart, K. Kim, V. Jampani, D. Fox, and J. Kautz, "Deepgmr: Learning latent gaussian mixture models for registration," in *ECCV*, 2020. 9, 13

[53] Z. J. Yew and G. H. Lee, "3dfeat-net: Weakly supervised local 3d features for point cloud registration," in *ECCV*, 2018. 14